arXiv:1604.02486v1 [cs.DM] 8 Apr 2016

The Salesman's Improved Paths: $3/2 + 1/34$ Integrality Gap and Approximation Ratio *

András Sebő[1] and Anke van Zuylen[2]

[1]CNRS, Univ. Grenoble Alpes, Optimisation Combinatoire (G-SCOP)
[2]Department of Mathematics, College of William & Mary, Williamsburg, VA, anke@wm.edu

April 12, 2016

Abstract

We give a new, strongly polynomial algorithm and improved analysis of the metric $s-t$ path TSP. The algorithm finds a tour of cost less than 1.53 times the optimum of the subtour elimination LP, while known examples show that 1.5 is a lower bound for the integrality gap.

A key new idea is the deletion of some edges, where the arising connectivity problems can be solved for a minor extra cost. On the one hand our algorithm and analysis extend previous tools, at the same time simplifying the framework. On the other hand new tools are introduced, such as a linear program that allows us to efficiently choose the probabilities of edges used for analyzing the reconnection cost and a new way of applying the matroid partition theorem and algorithm to different matroids, in order to write the subtour elimination optimum as a particular convex combination of spanning trees.

Furthermore, the 3/2 target-bound is easily reached in some relevant new cases.

keywords: path traveling salesman problem (TSP), approximation algorithm, Christofides' heuristic, the Chinese postman problem, matching theory, T-join, polyhedra

1 Introduction

In the Traveling Salesman Problem (TSP), we are given a set V of n "cities", a cost function $c : \binom{V}{2} \to \mathbb{Q}^{\geq 0}$, and the goal is to find a tour of minimum cost that starts and ends in the same city and visits each city exactly once. This "minimum length Hamiltonian circuit" problem is one of the most well-known problems of combinatorial optimization. In its full generality, the TSP is hopeless: it is not only *NP*-hard to solve but also to approximate, and even for quite particular lengths, since the Hamiltonian cycle problem in 3-regular graphs is *NP*-hard [15].

A condition on the cost function that helps in theory and is often satisfied in practice is known as the "triangle inequality" in complete graphs. A nonnegative function on the edges that satisfies this inequality is called a *metric* function. For a thoughtful and entertaining account of the difficulties and successes of the TSP, see Bill Cook's book [6].

If the cost function is a metric, then we may relax the problem into finding a Eulerian submultigraph (V, E^*) of minimum cost $\sum_{e \in E^*} c(e)$, since a traversal of this graph can be shortcut to a tour without increasing the cost. The multiplicities of edges can be restricted to be 0, 1 or 2.

*Research initiated during the Hausdorff Trimester Program in Combinatorial Optimization, Hausdorff Institute of Mathematics, Bonn; Partially supported by LabEx PERSYVAL-Lab (ANR 11-LABX-0025) and by a grant from the Simons Foundation (#359525, Anke Van Zuylen).

Christofides [5] gave a very simple $\frac{3}{2}$-approximation algorithm for the metric TSP: separate the problem into finding a minimum cost connected subgraph (minimum cost spanning tree) and completing this tree to an Eulerian subgraph, by correcting the parity of the constructed tree, S, i.e., by adding a minimum cost T_S-join, where T_S is the set of vertices that have odd degree in S (Chinese postman problem). Wolsey [25] observed that Christofides' algorithm also implies a bound on the strength of the well-known subtour elimination linear program (LP), that was introduced by Dantzig, Fulkerson and Johnson [8]:

$$\begin{aligned}
\text{Min} \qquad & c(x) := \sum_{e \in \binom{V}{2}} c(e)x(e) \\
\text{subject to:} \qquad & \sum_{e \in \delta(\{i\})} x(e) = f(\{i\}), \quad \text{for all } i \in V, && (1) \\
& \sum_{e \in \delta(U)} x(e) \geq f(U), \quad \text{for all } \emptyset \subsetneq U \subsetneq V, && (2) \\
& x(e) \geq 0, \quad \text{for all } e \in \binom{V}{2}. && (3)
\end{aligned}$$

where $f(U) \equiv 2$ for all sets U, and $\delta(U) = \{\{i,j\} \in E : i \in U, j \notin U\}$. The first set of constraints impose the condition that every vertex is visited once, and the second set of constraints, known as the subtour elimination constraints, ensure that every subset is visited at least once. Let x^* denote an optimal solution to the subtour elimination LP, and $OPT_{LP} := OPT_{LP}(c) := c(x^*)$ the optimal objective value. The integrality gap for metric TSP is the worst-case ratio, over all metric cost functions, of the cost of the optimal tour to OPT_{LP}. Wolsey [25] observed that x^* is in the the spanning tree polytope, and that $x^*/2$ is in the T-join polyhedron for any set T of even size, which implies that the integrality gap of the subtour elimination LP is at most $\frac{3}{2}$.

Despite significant effort, no improvements are known for either the approximation ratio or the integrality gap for the metric TSP.

A relevant generalization of the (metric) TSP is the (metric) $s-t$ path TSP. In this problem, the salesman starts in s, ends in t, and needs to visit every other city once. In other words, the goal is to find a Hamiltonian path from s to t of minimum cost. The subtour elimination LP for the $s-t$ path TSP is as above, but now defining $f(U) = 1$ if $|U \cap \{s,t\}| = 1$ and $f(U) = 2$ otherwise. Note that the special case $s = t$ is the regular traveling salesman problem.

In this paper, we describe a new, combinatorial, strongly polynomial algorithm and improved analysis with the new approximation and integrality ratio $3/2 + 1/34$ for the metric $s-t$ path traveling salesman problem. To achieve this, we also need *to delete edges from Christofides' trees*, in the hope that they will be automatically reconnected during parity correction; however, this is not always the case, and whenever it is not, we also need to invest separately in *reconnection*; for analyzing it, we use an LP for defining the distribution of a random choice for the reconnecting edges, we explore a new way of using matroid partition for finding the initial set of trees, and we consider a specific convex combination of T-joins rather than a minimum cost T-join.

To the best of our knowledge, the first relevant occurrence of possibly different s and t is in Hoogeveen's article [17], providing a Christofides-type approximation algorithm for the metric case, with an approximation ratio of 5/3 rather than 3/2. There had been no improvement until An, Kleinberg and Shmoys [1, 2] improved this ratio to $\frac{1+\sqrt{5}}{2} < 1.618034$ with a simple algorithm and an ingenious new framework for the analysis. The algorithm in [1, 2] is the *best-of-many* version of Christofides' algorithm. It first determines a minimum cost solution x^* of the subtour elimination LP. Then writing x^* as a convex combination of spanning trees and adding Christofides' *parity correction* for each, the algorithm outputs the best of the arising solutions. The best-of-many

algorithm was used by subsequent publications and is also used in the present work with some modification; mainly, we alter Christofides' algorithm by deleting certain edges from the spanning tree (see Section 2.1), and by then correcting both parity and connectivity.

To analyze the best-of-many algorithm, An, Kleinberg and Shmoys [1] come up with a certain "master formula" to bound the average parity correction cost, that was also used in subsequent publications. Sebő [20] improved the analysis further by separating each spanning tree in the convex combination into an $s-t$ path and its complement. Letting p^* be the indicator vector of the average of the $s-t$ paths, and q^* of the edges that are not on the $s-t$ path, we have that $x^* = p^* + q^*$. This separation makes it possible to use a simple and transparent master formula (see (6) in Section 2.2), but, surprisingly, until now the original, quite involved one was used by all the papers on $s-t$ path TSP following [1]. Moreover, p^* is further decomposed: it is the sum of naturally arising vectors (see x^Q later on). These vectors are getting an extended role in the present paper (Section 2.2).

For the $s-t$ path TSP on "graph metrics", that is, cost functions that are defined as the shortest path distances on a given unweighted graph, Gao [13] proves the existence of a tree in the support of LP solution x^* that has exactly one edge in so-called narrow cuts (see Section 2.1 below). For such a tree, it is possible to bound the cost of the parity correction by $c(x^*)/2$, as in Wolsey's version of Christofides' analysis. This allows Gao to give a very elegant proof of the approximation ratio 3/2 for graph metrics, a result that was first shown by [21] with a combinatorial, but more difficult proof.

Recently, Gottschalk and Vygen [16] significantly extended Gao's approach by proving the existence of a decomposition of x^* into certain "layers"; each layer consists of spanning trees that have exactly one edge in a subset of the narrow cuts. This will turn out to be needed in our analysis. We show that it is actually a *particular matroid partition*, which is also algorithmically useful for a strongly polynomial combinatorial algorithm (see Section 2.3 and Section 5). However, the constraints of the corresponding matroid partition (or basis packing) theorem of Edmonds [9] are apparently not straightforward to prove (see Appendix D for a weaker statement and the exposition of the problem); we therefore make use for the moment of [16] for checking the validity of Edmonds' matroid partition condition. The occurring particular matroid sums will probably remain a helpful tool for the s-t-path TSP problem, and we believe that a direct way of proving Edmonds' condition will eventually be found.

In the next section, we provide more details for the approaches described above in the form we need them in this paper, along with the novel algorithm and tools of our analysis.

2 Preliminaries, Algorithm and Outline of the Analysis

In this section we introduce our notation, terminology, and definitions. In the first subsection we also state our new algorithm. The second subsection gives an overview of the ideas we need to analyze our algorithm. The third subsection explains a particular phenomenon of the algorithm and its analysis: the necessity to choose a particular convex combination of trees.

2.1 The Algorithm: Best-of-Many with Deletion

Given a finite set V and costs $c : \binom{V}{2} \to \mathbb{Q}_{\geq 0}$ satisfying the triangle inequality, let x^* be the minimum cost solution to the subtour elimination LP for $s-t$ path TSP, and $E = \{e : x^*(e) > 0\}$. We will slightly abuse notation, and use the same notation for a (multi) subset of E and its own incidence vector in $\mathbb{Z}^E_{\geq 0}$. For a vector $z \in \mathbb{R}^E$ and any (multi)subset H of E, the usual notation $z(H) := \sum_{e \in H} z(e)$ is then just the scalar product of two vectors, as well as $c(z) := \sum_{e \in E} c(e)z(e)$.

We can write x^* as a convex combination of spanning trees; that is, there exist a collection of spanning trees $\mathcal{S}$ and coefficients $\lambda_S > 0$ for $S \in \mathcal{S}$ such that $\sum_{S\in\mathcal{S}} \lambda_S = 1$ and $x^* = \sum_{S\in\mathcal{S}} \lambda_S S$, and these can be found in strongly polynomial time with a combinatorial algorithm. (The existence follows from Tutte and Nash-Williams' theorems [23], [18]; the combinatorial algorithm from Cunningham [7]; see also Barahona [3] for a result specifically concerning spanning trees, and Schrijver [19, page 891] for an account of the best complexities so far.)

For a set of edges S, we let T_S be the set of odd degree vertices of S, i.e. those that have an odd number of incident edges from S. For two sets A, B, let $A\triangle B := (A \setminus B) \cup (B \setminus A)$ be the *symmetric difference* operation, which corresponds to the mod 2 sum of the incidence vectors. Then the vertex set with the "wrong" degree parity is $T_S\triangle\{s,t\}$. A T-join ($T \subseteq V$, $|T|$ is even) is a set of edges J such that $|J \cap \delta(\{i\})|$ is odd for $i \in T$ and $|J \cap \delta(\{i\})|$ is even for $i \in V \setminus T$. The operation "+" between sets means the disjoint union (sum of the multiplicity vectors). If S is a spanning tree and J_S a $T_S\triangle\{s,t\}$-join, $(V, S + J_S)$ is a spanning connected multigraph in which s and t have odd degree and every other vertex has even degree. We call the edge set of such a graph an $\{s,t\}$-tour.

The best-of-many Christofides' (BOMC) algorithm computes a minimum cost $T_S\triangle\{s,t\}$-join for every tree S with $\lambda_S > 0$, and outputs the resulting $\{s,t\}$-tour of minimum total cost. The cost of the BOMC solution is at most the cost of the $\{s,t\}$-tour we obtain by adding a $T_{\mathcal{S}}\triangle\{s,t\}$-join to a randomly chosen spanning tree $\mathcal{S}$ with $\Pr(\mathcal{S} = S) = \lambda_S$. Clearly, the expected cost of $\mathcal{S}$ is equal to $c(x^*)$; the main difficulty for the analysis of the s-t path TSP version of Christofides' algorithm is the "parity correction" part. It follows from Edmonds and Johnson [11] that the $T_S\triangle\{s,t\}$-join polyhedron is

$$\sum_{e\in\delta(U)} y(e) \geq 1 \text{ for all } U \text{ such that } |U \cap T_S\triangle\{s,t\}| \text{ is odd},$$
$$y(e) \geq 0 \text{ for all } e \in E.$$

As noted above, Wolsey [25] observed for the TSP (i.e. for $s = t$) and a solution x^* to the subtour elimination LP for the TSP that $x^*/2$ is in the T-join polyhedron for any set $T \subseteq V$, $|T|$ even. If $s \neq t$ this is not true anymore for a solution x^* to the corresponding subtour elimination LP since for cuts U separating s and t, called $s-t$ *cuts*, $f(U) = 1$ in the LP, and hence we are only guaranteed that $x^*(\delta(U))/2 \geq \frac{1}{2}$.

Following [1, 2], we say that $Q \subset E$ is *narrow* if $x^*(Q) < 2$. We let $\mathcal{Q}$ be the set of all narrow cuts, that is, $\mathcal{Q} = \{Q \subset E : x^*(Q) < 2\}$. Figure 1 shows an example of an optimal solution x^* to the subtour elimination LP for an $s-t$ path TSP and the narrow cuts $\mathcal{Q}$. An, Kleinberg and Shmoys [1, 2], observe that $\mathcal{Q}$ is a chain:

Lemma 1 (An, Kleinberg, Shmoys [1, 2]) *Let U_1, U_2 such that $s \in U_1 \cap U_2$, and $Q_i = \delta(U_i)$ is in $\mathcal{Q}$ for $i = 1, 2$. Then either $U_1 \subseteq U_2$, or $U_2 \subseteq U_1$.*

We include a brief proof (with some useful elements for the more general Lemma 15):

Proof: The lemma is a consequence of the submodular inequality that will be helpful in the following equality form:

$$x^*(\delta(A)) + x^*(\delta(B)) = x^*(\delta(A \cap B)) + x^*(\delta(A \cup B)) + 2x^*(A \setminus B, B \setminus A), \tag{4}$$

where (X, Y) is the set of edges with one endpoint in X, the other in Y. The equality is easily verified. Now to show Lemma 1 apply (4) to $A := U_1$, $B := V \setminus U_2$. If the lemma does not hold, $A \cap B$ and $A \cup B$ are both non-empty, strict subsets of V. Furthermore, note that $A \cap B$ contains neither of s, t and $A \cup B$ contains both of s, t. Hence, by constraint (2), the first and second terms in the right hand side of (4) are at least 2, contradicting that the left hand side is strictly smaller than 4 since $\delta(U_1), \delta(U_2) \in \mathcal{Q}$. □

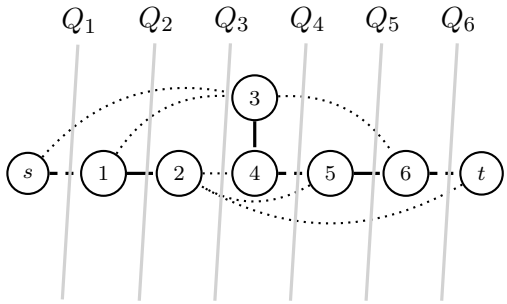

Figure 1: An example from Gao [14]: for full edges $x^*(e) = 1$, for dashed edges $x^*(e) = 2/3$ and for dotted edges $x^*(e) = 1/3$. The narrow cuts are indicated by gray lines and labeled $Q_1, \ldots, Q_6$.

An, Kleinberg and Shmoys [1, 2] observe that for a narrow cut $Q = \delta(U)$, the $T_S \triangle \{s,t\}$-join polyhedron has a constraint for U only if the spanning tree S has an even number of edges in Q. In Section 2.2, we give all necessary details of their approach and later works [20, 14, 24, 16] on analyzing the cost of the $T_S \triangle \{s,t\}$-join for $S \in \mathcal{S}$.

INPUT: Set V, two vertices, $s, t \in V$ and $c : \binom{V}{2} \to \mathbb{Q}^{\geq 0}$.
OUTPUT: An $\{s,t\}$-tour.
Let x^* be an optimal solution to the subtour elimination LP for the $s-t$ path TSP, and write x^* as a layered convex combination of $O(|E|)$ spanning trees (see the definition in Section 2.3, and a proof of existence in Section 5).
Let $\mathcal{S}$ be the collection of trees with positive coefficient.
for $S \in \mathcal{S}$ **do**
  Construct forest-based $\{s,t\}$-tour:
  Let $L(S)$ be the set of lonely edges of S, and $F := F(S) := S \setminus L(S)$.
  Let y_F be the vector in the $T_F \triangle \{s,t\}$-join polyhedron defined in Section 3.
  Write y_F as a convex combination of at most $O(|E|)$ different $T_F \triangle \{s,t\}$-joins. Let $\mathcal{J}_F$ be the collection of $T_F \triangle \{s,t\}$-joins with positive coefficient.
  for $J_F \in \mathcal{J}_F$ **do**
    Contract the components of $(V, F + J_F)$, and let $2D_{F+J_F}$ be the edge set of a doubled minimum cost spanning tree on the contracted graph.
    $P_1(F, J_F) = F + J_F + 2D_{F+J_F}$.
  end for
  Let $P_1(S)$ be a minimum cost $\{s,t\}$-tour among $\{P_1(F, J_F)\}_{J_F \in \mathcal{J}_F}$.
  Construct Christofides-based $\{s,t\}$-tour:
  $P_2(S) = S + J_S$, where J_S is a minimum cost $T_S \triangle \{s,t\}$-join.
end for
Return a minimum cost $\{s,t\}$-tour from among $\bigcup_{S \in \mathcal{S}} \{P_1(S), P_2(S)\}$.

Algorithm 1: Best-of-Many With Deletion (BOMD)

In this paper, we modify the best-of-many Christofides' algorithm by *deleting* certain edges called *lonely* from the spanning trees $S \in \mathcal{S}$. We denote by $L(S)$ the set of lonely edges of tree S, where we require for all $e \in L(S)$ that the two components of $S \setminus \{e\}$ are the two sides of a narrow cut, that is,

$$L(S) \subseteq \big\{e \in S : \text{there exists } Q \in \mathcal{Q} \text{ such that } Q \cap S = \{e\}\big\}. \tag{5}$$

Note that we do not require equality in (5), and this flexibility will enable us to choose the lonely

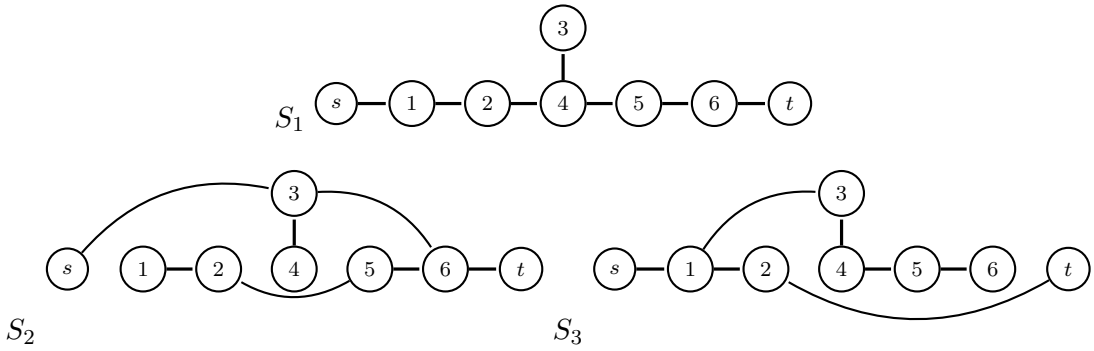

Figure 2: x^* from Figure 1 expressed as a layered convex combination of three spanning trees S_1, S_2, S_3, each with multiplier $\lambda_{S_i} = 1/3$. Note that S_1 is a so-called Gao-tree (see Section 2.3). All edges but $\{3, 4\}$ are lonely in S_1, and in S_2, S_3 only the edges incident to s and t are lonely.

edges to satisfy additional conditions. For each tree $S \in \mathcal{S}$, we delete its lonely edges to obtain a forest $F(S) := S \setminus L(S)$. We write F instead of $F(S)$ whenever S is clear from the context. We then add a $T_F \triangle \{s, t\}$-join to obtain a graph in which each vertex except for s and t has even degree, and finally, we add a doubled spanning tree on the components of this graph to get an $\{s, t\}$-tour. Note that adding a spanning tree on the components suffices to make the graph connected, and we double it in order to keep the degree parities.

In order to bound the cost of the doubled spanning tree, we will need to use properties of the $T_F \triangle \{s, t\}$-join; more precisely, we will take a fractional solution in the $T_F \triangle \{s, t\}$-join polyhedron and write it as a convex combination of $T_F \triangle \{s, t\}$-joins, and then repeat the computation for each of the $T_F \triangle \{s, t\}$-joins in the convex combination. In addition to computing several forest-based $\{s, t\}$-tours (with different $T_F \triangle \{s, t\}$-joins) for each $S \in \mathcal{S}$, the algorithm also computes a Christofides-based $\{s, t\}$-tour, and returns the least expensive among all constructed $\{s, t\}$-tours. (See Algorithm 1 for a complete formalized description.)

2.2 The Analysis and its Setting: Deletion, Correction, Completion, Reconnection, Balancing

Given $x^* = \sum_{S \in \mathcal{S}} \lambda_S S$, where $\mathcal{S}$ is a set of spanning trees, $\lambda_S > 0$ for $S \in \mathcal{S}$, and $\sum_{S \in \mathcal{S}} \lambda_S = 1$, recall that this coefficient vector can be interpreted as a probability distribution. We will slightly abuse notation and let $\mathcal{S}$ denote both the support of the convex combination, and also a tree-valued random variable with $\Pr(\mathcal{S} = S) = \lambda_S$. (See Figure 2 for such a convex combination, where x^* is from Figure 1.)

We will also take distributions over convex combinations of T-joins. We let $\mathbb{E}[\cdot]$ denote the expectation or average of a random vector in $\mathbb{R}^E$; for example $\mathbb{E}[\mathcal{S}] = x^*$.

For a spanning tree S, let $S(s, t)$ be the $s - t$ path of S. Various notations for trees S will be inherited by $\mathcal{S}$ (and similarly for T-joins): for instance, $\mathcal{S}(s, t)$ is the path-valued random variable which is equal to $S(s, t)$ on the event $\mathcal{S} = S$.

2.2.1 Bounding the cost of the Christofides-based $\{s, t\}$-tours

To analyze the average cost of the Christofides-based $\{s, t\}$-tours, we detail an observation of [20] used in the analysis of the best-of-many Christofides' algorithm.

Observation 2 *Since $S = S(s, t) + S \setminus S(s, t)$, and for any $S \in \mathcal{S}$, $S + S \setminus S(s, t)$ is an $\{s, t\}$-tour, defining $p^* := \mathbb{E}[\mathcal{S}(s, t)]$, $q^* := \mathbb{E}[\mathcal{S} \setminus \mathcal{S}(s, t)]$, $x^* = \mathbb{E}[\mathcal{S}] = p^* + q^* \in \mathbb{R}^E$, the cost of the solution*

found by the best-of-many Christofides' algorithm can be bounded by

$$\mathbb{E}\left[c(\mathcal{S}) + c(\mathcal{S} \setminus \mathcal{S}(s,t))\right] = c(x^*) + c(q^*) = 2c(x^*) - c(p^*).$$

When $c(q^*)$ is sufficiently small compared to $c(p^*)$, then this bound provides the best approximation guarantee; when it is large, another bound is better, and it is shown in [20] that the best of the two provides a bound of 8/5 on the approximation and integrality ratio. In Appendix A, we rederive this bound in a surprisingly simple way.

In fact, since we will be taking the minimum of the Christofides-based $\{s,t\}$-tour and the forest-based $\{s,t\}$-tour, this observation allows us to "erase" $\mathcal{S} \setminus \mathcal{S}(s,t)$ from further consideration for parity correction, i.e., for analyzing the cost of the forest-based $\{s,t\}$-tours, we will only take the complementary $\mathcal{S}(s,t)$ into account.

2.2.2 Basic parity correction

Concretely, to analyze the parity correction for the forests, we will always use

$$\frac{x^*}{2} + \gamma\, S(s,t), \tag{6}$$

as a so-called *basic parity correction*, where γ is a chosen parameter between 0 and 1/2. The basic parity correction plays a similar role to the vector $\beta\frac{x^*}{2} + (1-\beta)S$ that was used in the analysis of [1], which was also adopted by subsequent work [20, 14, 24, 16].

For $\gamma = 1/2$ the basic parity correction $\frac{x^*}{2} + \frac{1}{2}S(s,t)$ for any $S \in \mathcal{S}$ gives a feasible solution to the T-join polyhedron for any set T of even size, so in particular it allows us to bound the cost of the $T_S\triangle\{s,t\}$-join. On average, the basic parity correction with $\gamma = 1/2$ will equal $\frac{x^*}{2} + \frac{p^*}{2}$, and combined with Observation 2, we get the integrality gap and approximation ratio $\min\left\{\frac{3}{2}c(x^*) + \frac{1}{2}c(p^*), 2c(x^*) - c(p^*)\right\} \le \frac{5}{3}c(x^*)$ equal to Hoogeveen's ratio. We can similarly use $\frac{x^*}{2} + \frac{1}{2}S(s,t)$ as a fractional $T_F\triangle\{s,t\}$-join for the forest $F = S \setminus L(S)$; however, the summand $\frac{1}{2}\mathcal{S}(s,t)$ is too expensive to get a good bound.

To improve the ratio further, we will need $\gamma < 1/2$, but then the basic parity correction vector may not be feasible for the $T_F\triangle\{s,t\}$-join polyhedron. Hence, as in previous analyses of the best-of-many algorithm, we will need to add a *parity completion*[1] vector to the basic parity correction vector so as to ensure the constraints of the $T_F\triangle\{s,t\}$-join polyhedron are satisfied for the narrow cuts Q with $x^*(Q) < 2 - 2\gamma$, in which F contains an even number of edges.

However, we cannot reduce γ to 0, since the completion vectors are too expensive for large narrow cuts. The treshold $2 - \gamma$ for being "large" will be chosen optimally at the end of Section 3.

2.2.3 Parity completion

The parity completion vector will add (fractional) edges for each narrow cut Q if $|Q \cap F|$ is even. If we define the set of lonely edges in tree S by taking (5) with equality, the expectation of the incidence vectors of the lonely edges for each narrow cut Q is equal to a vector defined in Sebő [20], which was used in [20] for defining good parity completion vectors. In particular, for each narrow cut Q, we introduce a vector $x^Q \in \mathbb{R}^E$, where

$$x^Q(e) := \Pr(e \text{ is the lonely edge of } \mathcal{S} \text{ in } Q). \tag{7}$$

[1] Our terminology differs from that in [1], where the term "correction vector" is used for what we call the "parity completion vector". We use the term "parity correction vector" for a vector that is in the $T_F\triangle\{s,t\}$-join polyhedron, and it will be composed of basic parity correction and parity completions.

If (5) is satisfied with equality, then $\Pr(|Q \cap \mathcal{S}| = 1) = x^Q(Q)$ and

$$x^Q(Q) \geq 2 - x^*(Q), \tag{8}$$

because $x^*(Q) = \mathbb{E}[|Q \cap \mathcal{S}|] = \sum_{k=1}^{\infty} k \Pr(|Q \cap \mathcal{S}| = k) \geq \Pr(|Q \cap \mathcal{S}| = 1) + 2(1 - \Pr(|Q \cap \mathcal{S}| = 1))$. This is also observed in equation (3) of [2] as Markov's inequality.

As mentioned after (5), $L(S)$ is allowed to be a proper subset of the right hand side of (5), and the lonely edges will be chosen to satisfy additional properties. One property that we require now from the set of lonely edges $L(S)$ in addition to (5), is that

$$x^Q(Q) = 2 - x^*(Q), \tag{9}$$

where we recall that $x^Q(Q) = \sum_{e \in Q} \Pr(L(\mathcal{S}) \cap Q = \{e\})$. This equality can clearly be satisfied by removing edges from $L(S)$ for some $S \in \mathcal{S}$ appropriately.

Hence, we have that $\frac{x^*(Q)+x^Q(Q)}{2} = 1$, on every cut. In other words, adding $x^Q/2$ to the basic parity correction vector for every Q such that $|Q \cap F|$ is even gives a suitable vector to *complete* parity correction, that is, for *parity completion*.

On the other hand, if we let $\mathcal{F}$ be the random forest obtained by deleting $L(\mathcal{S})$ from $\mathcal{S}$, then we also have that $x^Q(e)$ is the *deletion probability* for edge e and cut Q. So by deleting the lonely edges from $\mathcal{S}$, we "save up" for parity completion. In Section 3, we show how to put these ingredients together to construct an appropriate parity correction vector y_F, and we analyze its mean cost over $S \in \mathcal{S}$.

2.2.4 Reconnection

Finally, we need to bound the average cost of the *doubled minimum cost spanning tree* $2D_{F+J_F}$ *between the components of* $(V, F + J_F)$, meaning the doubled edge-set of a spanning tree in the graph whose vertex-set is the set of contracted components of $(V, F + J_F)$. Note that doubling a minimum cost spanning tree between the components of $(V, F + J_F)$ gives the minimum possible cost of the reconnection of $F + J_F$ if using *only* doubled edges; for bounding the expectation of $2D_{\mathcal{F}+J_{\mathcal{F}}}$ we can therefore consider any set of doubled edges that reconnect the forest.

For a tree $S \in \mathcal{S}$, denote by $\mathcal{Q}(S) \subseteq \mathcal{Q}$ the set of *lonely cuts* of S: cuts Q for which $|Q \cap S| = 1$ for which the unique edge in $Q \cap S$ is in $L(S)$. For $Q \in \mathcal{Q}(S)$, we denote by e_S^Q the (incidence vector of) the unique edge of $Q \cap S$.

Since $F = S \setminus L(S)$, J_F must contain at least one edge in every cut in $Q \in \mathcal{Q}(S)$, since $|F \cap Q|$ is even (and equal to 0). Now, if each edge in J_F is contained in at most one lonely cut, then $(V, F + J_F)$ is connected, and the doubled spanning tree is empty, its cost is 0. We will call the edges that are contained in more than one lonely cut of S the *bad edges* for S, and denote them by $B(S)$, that is,

$$B(S) = \{e \in E : \text{there exist } Q_1 \neq Q_2 \in \mathcal{Q}(S),\ e \in Q_1 \cap Q_2\}.$$

Ideally, we would like a low cost $T_F \triangle \{s,t\}$-join J_F (or, a fractional $T_F \triangle \{s,t\}$-join y_F) that has no edges in $B(S)$. Figure 3 shows a forest F and a $T_F \triangle \{s,t\}$-join J_F with two bad edges.

As we sketched above, our fractional $T_F \triangle \{s,t\}$-join is going to consist of the basic parity correction $\frac{x^*}{2} + \gamma S(s,t)$ plus a parity completion vector for which we will use the x^Q vectors for $Q \in \mathcal{Q}$. There is no way to avoid that the basic parity correction vector contains bad edges, but we will choose our convex combination of x^* into spanning trees in such a way that parity completion does not add more bad edges (see sections 2.3 and 5 for more details).

Given the parity correction vector y_F that we construct, we consider a random $T_F \triangle \{s,t\}$-join $\mathcal{J}_F$ with $\mathbb{E}[\mathcal{J}_F] = y_F$. In Section 4 we give an upper bound on the expected cost of the doubled spanning tree $2D_{F+\mathcal{J}_F}$, where the expectation is over $\mathcal{J}_F$. For any $e \in E$, we define $\mathcal{Q}(S,e)$ as the

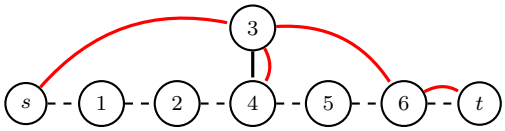

Figure 3: The forest $F_1 := F(S_1)$ obtained by removing the lonely edges from S_1 in Figure 2 contains only the edge $\{3,4\}$ indicated in (solid) black, the dashed edges are the lonely edges of S_1. The (solid) red edges give a possible $T_{F_1} \triangle \{s,t\}$-join J_{F_1}, where edges $\{s,3\}$ and $\{3,6\}$ are bad edges. An upper bound on $2D_{F_1+J_{F_1}}$ can be obtained by adding two copies of two out of the three lonely edges $\{s,1\},\{1,2\},\{2,4\}$ and adding two copies of either $\{4,5\}$ or $\{5,6\}$.

set of lonely cuts of S that contain e; note that $b \in B(S)$ if and only if $|\mathcal{Q}(S,b)| \geq 2$. We can upper bound the cost of $2D_{F+\mathcal{J}_F}$ by, for each bad edge $b \in \mathcal{J}_F$, taking two copies of each lonely edge in $S \cap \mathcal{Q}(S,b)$ except possibly one such lonely edge; so the cost can be upper bounded by doubling the cost of

$$\bigcup_{b \in \mathcal{J}_F} \{e_S^Q : \text{ for all but at most one } Q \text{ in } \mathcal{Q}(S,b)\}. \tag{10}$$

The challenge that is addressed in Section 4 is to decide for which $Q \in \mathcal{Q}(S,b)$ we do *not* add two copies of the edge e_S^Q. The goal is to choose it so that on average (over all $J_F \in \mathcal{J}_F$), we get a good bound on the number of times we use e_S^Q, and hence on the cost of $2D_{F+\mathcal{J}_F}$. Note that $\mathbb{E}\left[e_S^Q\right] = x^Q$, and we thus express the cost of reconnection in terms of the vectors $\{x^Q\}_{Q \in \mathcal{Q}}$ as well.

2.2.5 Balancing

To bound the average cost of the forest-based $\{s,t\}$-tours, we now simply add up the average cost of the forest plus parity correction (consisting of basic parity correction and parity completion) plus reconnection. By noting that $\sum_{Q \in \mathcal{Q}} x^Q \leq \mathbb{E}[\mathcal{S}(s,t)] = p^*$ (since the lonely edges of spanning trees S must be contained in $S(s,t)$), all terms of the bound on the average cost of the forest-based $\{s,t\}$-tours will be expressable in terms of x^* and p^*. Finally, the analyis can take the minimum of the Christofides-based $\{s,t\}$-tours (which by Observation 2 has cost decreasing in $c(p^*)$) and the cost of the forest-based $\{s,t\}$-tours is increasing in $c(p^*)$. The worst case for the minimum of the two is when the two are *balanced*, that is equal.

2.3 Matroid Partition to Avoid Additional Reconnection

Besides the improved algorithm, and the main ideas of the analysis, there is an important, technical looking detail that has to be handled: as explained in the previous subsection, not only parity, but connectivity also has to be corrected if the $T_F \triangle \{s,t\}$-join contains bad edges. We show in Section 5 that when choosing a convex combination of trees in a particular way that can be computed with Edmonds' "matroid partition algorithm", parity completion does not contribute to the probability of bad edges occurring in the $T_F \triangle \{s,t\}$-join. Let us introduce the reader to this nontrivial matter.

In a beautiful paper [13], Gao showed that there always exists a spanning tree S in the support of x^* (that is, in $S \subseteq \{e : x^*(e) > 0\}$) such that $|S \cap Q| = 1$ for every narrow cut Q. Whenever there exists such a minimum cost spanning tree that satisfies this property (like for "graph-metrics") the 3/2 ratio and integrality gap follow straightforwardly. A tree in which every narrow cut is *lonely*, that is, contains a single edge of the tree, is called a *Gao-tree*.

The interesting idea of *choosing the convex combination of spanning trees in a particular way*, so as to decrease the cost of parity completion, was introduced in Vygen [24]. The claimed progress in the ratio was only 0.001, but the idea of reassambling the trees $S \in \mathcal{S}$ which participate in the convex

combination is further developed by Gottschalk and Vygen [16]. They generalize the concept of Gao-trees, and the reassambling leads to a powerful result: a convex combination using generalized Gao-trees. Although the use of local changes is constructive, the algorithm it implies is not a polynomial time algorithm. However, using an approximate solution to this convex combination that can be found in polynomial time, Gottschalk and Vygen state that the best-of-many algorithm gives a 1.566-approximation and integrality ratio for the $s-t$ path TSP.

In Section 5 we prove that (generalized) Gao-trees form a matroid, and that a convex combination into these trees can be found as a corollary of Edmonds' matroid partition algorithm [10] (see also [12, Theorem 13.3.1], [19, Corollary 42.1a]). As a consequence, a "suitable" convex combination is constructed directly (without a posteriori reassambling or exchanges), in strongly polynomial time, from the matroid partition algorithm applied to the matroids of (generalized) Gao-trees.

We note that, independent of our work, Kanstantsin Pashkovich also pointed out that this convex combination can be found (in weakly polynomial time) using the ellipsoid method [Jens Vygen, private communication, 2016].

We will use the term "layered" to refer to the resulting convex combination: We say that the spanning tree S of G with lonely edges $L(S)$ and consequently defined lonely cuts $\mathcal{Q}(S)$ is *layered*, if

$$\text{for each } Q \in \mathcal{Q}(S) \text{ and cut } Q' \text{ such that } x^*(Q') \leq x^*(Q) \text{ we have } Q' \in \mathcal{Q}(S),$$

that is, Q' is also lonely in S if it is smaller than another cut Q lonely in S. We say that the tree-valued random variable $\mathcal{S}$, or the convex combination of spanning trees $\mathcal{S}$ with coefficients λ_S for $S \in \mathcal{S}$ is *layered*, if $\mathbb{E}[\mathcal{S}] = x^*$, (5), (9) hold, and each $S \in \mathcal{S}$ is layered.

In Section 5 the existence and (strongly) polynomial algorithm for finding a layered convex combination will be confirmed. We can view these results as strengthenings of (9), which states that the larger the cut, the more rarely it is lonely; in a layered convex combination this is true in the stronger sense of the containment (implication) of events. In order to achieve this, we exploit the possibility of proper containment in (5), i.e., $\mathcal{Q}(S)$ may be a proper subset of $\{Q \in \mathcal{Q} : |Q \cap S| = 1\}$.

We finish this section by showing how layered convex combinations allow us to avoid additional reconnection costs. We also state the theorem that converts this to a simple bound on the total of reconnection costs; the proof is deferred to Section 4.

Lemma 3 *Let $\mathcal{S}$ be layered, $S \in \mathcal{S}$, and $Q \in \mathcal{Q} \setminus \mathcal{Q}(S)$. Then $x^Q(B(S)) = 0$.*

Proof: We prove $x^Q(Q') = 0$ for each $Q' \in \mathcal{Q}(S)$. Indeed, since S is layered, $x^*(Q') < x^*(Q)$, for otherwise $Q \in \mathcal{Q}(S)$. But then using that $\mathcal{S}$ is layered, Q' is lonely in every tree $S' \in \mathcal{S}$ where Q is lonely, so Q and Q' cannot have a common edge in such a tree S', proving $x^Q(Q') = 0$.

Therefore, the edges e with $x^Q(e) > 0$ are not contained in any $Q \in \mathcal{Q}(S)$ at all, whereas the edges of $B(S)$ are contained in at least two cuts of $\mathcal{Q}(S)$. □

Algorithm 1 loops through a layered $\mathcal{S}$, and has an inner loop through the support $\mathcal{J}_F$ of a convex combination of $T_F \triangle \{s,t\}$-joins, for each $F \in \mathcal{F} := \mathcal{S} \setminus L(\mathcal{S})$. Consider the corresponding $T_F \triangle \{s,t\}$-join valued random variable $\mathcal{J}_F$ (with $\mathbb{E}[\mathcal{J}_F] = y_F$) that *accompanies* F. (In Algorithm 1, the inner loop determines a minimum cost reconnecting edge-set $2D_{F+J_F}$ for every J_F in the support of $\mathcal{J}_F$, thus derandomizing the random $T_F \triangle \{s,t\}$-join $\mathcal{J}_F$ that accompanies F.)

Theorem 4 *Assume $\mathcal{S}$ is layered, $S \in \mathcal{S}$, and $F = S \setminus L(S)$ is accompanied by $\mathcal{J}_F$. Then*

$$\mathbb{E}[c(2D_{F+\mathcal{J}_F})] \leq \sum_{Q \in \mathcal{Q}(S)} (x^*(Q) - 1) c(e_S^Q).$$

3 The New Ratio

In this section we prove the main result, relying on Theorem 4. It may be useful for the reader to consult Appendix B at this point, where we show how the bound 3/2 can be proved more simply for special cases, without the need for Theorem 4.

Theorem 5 *The Best-of-Many With Deletion (BOMD) algorithm returns a solution to the $s-t$ path TSP of cost at most $\left(\frac{3}{2}+\frac{1}{34}\right) OPT_{LP}$.*

Proof: We analyze the different parts of the forest-based $\{s,t\}$-tour $P_1(S)$ averaged over $S \in \mathcal{S}$.

Forest: Let $\mathcal{F} := \mathcal{S} \setminus L(\mathcal{S})$.

Expected Value of Forest : $\mathbb{E}[\mathcal{F}] = x^* - \sum_{Q \in \mathcal{Q}} x^Q$.

Parity Correction (PC): For each forest F, we define y_F as the sum of the following vectors:

Basic Parity Correction (BP): $\frac{1}{2}x^* + \gamma\, S(s,t)$.
Note that (BP) is enough to ensure that y_F satisfies the constraints of the $T_F \triangle \{s,t\}$-join polyhedron for all $Q = \delta(U)$ with $x^*(Q) \geq 2 - 2\gamma$.
Expected Value of (BP): $\dfrac{1}{2}x^* + \gamma\, p^*$.

Parity Completion for Empty Cuts (PCE): $\left(1 - \dfrac{1}{2}x^*(Q) - \gamma\right) e_S^Q$.

We recall that e_S^Q is the (indicator vector of) the lonely edge of S in Q. We add this (PCE) vector to y_F for every $Q \in \mathcal{Q}(S)$ such that $x^*(Q) \leq 2-2\gamma$. Note that then (BP)+(PCE) suffices to ensure that y_F satisfies the constraint of the $T_F \triangle \{s,t\}$-join polyhedron for each narrow cut Q such that $Q \cap F = \emptyset$. Since $\mathbb{E}\left[e_{\mathcal{S}}^Q\right] = x^Q$, we have:

Expected Value of (PCE): $\displaystyle\sum_{Q \in \mathcal{Q}:\, x^*(Q) \leq 2-2\gamma} \left(1 - \frac{1}{2}x^*(Q) - \gamma\right) x^Q$.

Parity Completion for Non-Empty Even Cuts (PCL): $\left(\dfrac{1 - \frac{1}{2}x^*(Q) - \gamma}{2 - x^*(Q)}\right) x^Q$.

We add this vector to y_F for every Q such that $x^*(Q) \leq 2-2\gamma$ and $|F \cap Q|$ is even and at least 2. By (9), y_F satisfies the constraints of the $T_F \triangle \{s,t\}$-join polyhedron also for narrow cuts such that $|F \cap Q|$ is even and non-zero.
Contrary to (PCE), (PCL) does not depend on (the outcome of the random variable) $\mathcal{S}$, but only on the cut $Q \in \mathcal{Q}$. However, the choice of adding it or not, depends on $\mathcal{S}$: we add it only if $|\mathcal{F} \cap Q|$ is even and at least two, the probability of which can be bounded by $\Pr(|\mathcal{S} \cap Q| \geq 2)$, which is at most $1 - (2 - x^*(Q))$ by (9), so we add (PCL) with probability at most $x^*(Q) - 1$:

Expected Value of (PCL): $\displaystyle\sum_{Q \in \mathcal{Q}:\, x^*(Q) \leq 2-2\gamma} (x^*(Q) - 1) \left(\frac{1 - \frac{1}{2}x^*(Q) - \gamma}{2 - x^*(Q)}\right) x^Q$.

The combination of the parts (BP) + (PCE) + (PCL) gives a fractional solution to the $T_F \triangle \{s,t\}$-join polyhedron for any fixed tree S and its uniquely determined forest $F = S \setminus L(S)$:

$$\begin{aligned} y_F \;=\; & \frac{1}{2}x^* + \gamma\, S(s,t) + \sum_{Q \in \mathcal{Q}:\, x^*(Q) < 2-2\gamma,\, |Q \cap F| = 0} \left(1 - \frac{1}{2}x^*(Q) - \gamma\right) e_S^Q \\ & + \sum_{Q \in \mathcal{Q}:\, x^*(Q) < 2-2\gamma,\, |Q \cap F| \geq 2, \text{even}} \left(\frac{1 - \frac{1}{2}x^*(Q) - \gamma}{2 - x^*(Q)}\right) x^Q. \end{aligned} \tag{11}$$

Reconnection: For each $S \in \mathcal{S}$, define the random $T_F \triangle \{s,t\}$-join $\mathcal{J}_F$ with $\mathbb{E}[\mathcal{J}_F] = y_F$. By Theorem 4: *Expected Value of Reconnection Cost:* $\mathbb{E}[c(2D_{\mathcal{F}+\mathcal{J}_F})] \leq \sum_{Q\in\mathcal{Q}}(x^*(Q)-1)c(x^Q)$.

TOTAL: If we add up the costs of the bounds on the different parts of the solution, we get $\frac{3}{2}c(x^*) + \gamma c(p^*)$ plus the sum over all $Q \in \mathcal{Q}$ of some multiple (that depends on $x^*(Q)$) of $c(x^Q)$. In particular, if $2-2\gamma < x^*(Q) < 2$, the multiplier of $c(x^Q)$ is $-1 + x^*(Q) - 1 = x^*(Q) - 2 \leq 0$, and if $x^*(Q) \leq 2-2\gamma$ it is

$$-1 + \left(1 - \frac{1}{2}x^*(Q) - \gamma\right)\left(1 + \frac{x^*(Q)-1}{2-x^*(Q)}\right) + x^*(Q) - 1 = \frac{2 - x^*(Q) - 2\gamma}{2(2-x^*(Q))} + x^*(Q) - 2.$$

We choose γ so that this multiplier is nonpositive for all $Q \in \mathcal{Q}$, that is, we want

$$2 - x^*(Q) - 2\gamma + 2(2-x^*(Q))(x^*(Q)-2) \leq 0 \text{ that is, } (2 - x^*(Q) - \frac{1}{4})^2 + \gamma - \frac{1}{16} \geq 0.$$

The minimum of γ for which this is satisfied for all $1 \leq x^*(Q) \leq 2$ is $\gamma = \frac{1}{16}$, where we note that equality holds if and only if $x^*(Q) = \frac{7}{4}$.

We have thus bounded the average cost of the $P_1(\mathcal{S})$ solutions by $\frac{3}{2}c(x^*) + \frac{1}{16}c(p^*)$. By Observation 2, the average cost of the $P_2(\mathcal{S})$ solutions is at most $2c(x^*) - c(p^*)$, and therefore the cost of the solution returned by (BOMD) is at most

$$\min\left\{\frac{3}{2}c(x^*) + \frac{1}{16}c(p^*), 2c(x^*) - c(p^*)\right\} = \left(\frac{3}{2} + \frac{1}{34}\right)c(x^*).$$

□

4 Bounding the Reconnection Cost

In this section we provide a bound, tight in some sense, on the cost of the doubled spanning tree $c(2D_{F+J_F})$ between the components of $(V, F + J_F)$, as explained in Section 2.2.4.

We fix S and $F = S \setminus L(S)$, and in this section, we think of $\mathcal{J}_F$ as representing a random $T_F \triangle \{s,t\}$-join, with $\mathbb{E}[\mathcal{J}_F] = y_F$. Recall that (10) gives us a way to upper bound the expected cost of $2D_{F+\mathcal{J}_F}$: for each bad edge $b \in \mathcal{J}_F$, we add two copies of the lonely edges e_S^Q for *all but at most one of the lonely cuts* $Q \in \mathcal{Q}(S,b)$ *that contain* b. We will say that the edges for which two copies are added are *used for reconnection*. Obviously, the best bound is obtained by choosing to omit the most expensive lonely edge; however, for our analysis, we will randomly choose the cut $Q \in \mathcal{Q}(S,b)$ for which we *do not* add two copies of the edge e_S^Q for reconnection. Note that we are free to determine the probability distribution and, for simplicity, we also allow not choosing any of the cuts. We denote

$$x(b,Q) := \Pr(e_S^Q \text{ is not used for reconnection} \mid b \in \mathcal{J}_F),$$

and describe a system of linear inequalities that turns out to express that $\{x(b,Q) : b \in B(S), Q \in \mathcal{Q}(S)\}$ is a suitable set of conditional probabilities (as will be shown in Lemma 6):

$$(D0) \qquad x(b,Q) \geq 0, \quad \text{for all } b \in B(S), Q \in \mathcal{Q}(S,b),$$

$$(D1) \qquad \sum_{Q\in\mathcal{Q}(S,b)} x(b,Q) \leq 1 \text{ for all } b \in B(S),$$

$$(D2) \qquad \sum_{b \in B(S) \cap Q} x^*(b)x(b,Q) \geq r_Q, \text{for all } Q \in \mathcal{Q}(S),$$

where $r_Q := 1 - x^*(Q \setminus B(S))$.

Note that r_Q can be negative, in which case the constraint is vacuously satisfied.

The conclusion of the following lemma, Lemma 6, will lead to Theorem 4; the conditions of Lemma 6 will be assured by two additional lemmas: Lemma 7 and Lemma 3.

Lemma 6 *Let $Q \in \mathcal{Q}(S)$, and suppose $\Pr(b \in \mathcal{J}_F) \leq x^*(b)/2$ for each $b \in B(S)$. Let $x(b,Q) := \Pr(e_S^Q$ is not used for reconnection $|b \in \mathcal{J}_F)$, then*

$$\Pr(e_S^Q \text{ is used for reconnection }) \leq \frac{x^*(Q) - 1}{2},$$

provided $\{x(b,Q) : Q \in \mathcal{Q}(S), b \in B(S)\}$ satisfy (D2).

Proof: By the law of total probability:

$$\Pr(e_S^Q \text{ is used for reconnection }) = \sum_{b \in B(S)} (1 - x(b,Q)) \Pr(b \in \mathcal{J}_F),$$

where we can substitute the condition $\Pr(b \in \mathcal{J}_F) \leq x^*(b)/2$, and arrive at

$$\Pr(e_S^Q \text{ is used for reconnection }) \leq \sum_{b \in B(S)} \frac{x^*(b)}{2} (1 - x(b,Q)), \text{ for all } Q \in \mathcal{Q}(S).$$

Therefore, the claimed inequality holds provided that

$$\sum_{b \in B(S)} (x^*(b)/2)(1 - x(b,Q)) \leq \frac{x^*(Q) - 1}{2}, \text{ for all } Q \in \mathcal{Q}(S).$$

Multiplying by 2 and rearranging, we get that this inequality is equivalent to

$$-\sum_{b \in B(S)} x^*(b)x(b,Q) \leq x^*(Q) - 1 - \sum_{b \in B(S)} x^*(b), \text{ for all } Q \in \mathcal{Q}(S),$$

which is exactly $(D2)$ multiplied by -1. □

Despite the seemingly rough bounding, all of the applied inequalities in the series are best possible in the following sense: the resulting inequality cannot be improved with our methods, not even if there is only one narrow-cut-size in addition to 1 (which is a special case of Appendix B.2).

Lemma 7 *The system of linear inequalities (D0), (D1), (D2) is feasible for all $S \in \mathcal{S}$.*

Proof: Let $S \in \mathcal{S}$. Farkas' lemma provides a necessary and sufficient condition for feasibility. The special structure of the bipartite matching constraints $(D1)$, $(D2)$ makes possible to simplify this condition to the following Kőnig-Hall type characterization (in fact equivalent to the Kőnig-Hall theorem for f-factors or transportation problems:

Claim 1 *The system of linear inequalities (D0), (D1), (D2) is feasible if and only if for all $\mathcal{Q}' \subseteq \mathcal{Q}(S)$:*

$$\sum_{Q \in \mathcal{Q}'} r_Q \leq x^*(\cup_{Q \in \mathcal{Q}'} Q \cap B(S))$$

The proof that this condition is indeed necessary and sufficient can be done for instance by a reduction to a transportation problem. Let K be the least common denominator of x^*. We create a transportation problem with $|\mathcal{Q}(S)|$ demand nodes, one for each $Q \in \mathcal{Q}(S)$, with demand r_Q, and $K|B(S)|$ supply nodes, each with supply 1. Note that there are thus K supply nodes $b^1, \ldots, b^K$ corresponding to bad edge b; for each $Q \in \mathcal{Q}(S)$ such that $b \in Q$, we create K arcs (b^i, Q) for $i = 1, \ldots, K$. One can show that if there exists a solution to this transportation problem then there exists a solution in which for each b, Q, the flow on the arcs (b^i, Q) is equal for $i = 1, \ldots, K$. Such a solution is exactly a solution to $(D0)$, $(D1)$, $(D2)$.

For self-containedness and since a formal reduction would take at least as much space, we provide a direct proof in terms of inequalities in Appendix C.

Therefore, what we have to prove is that the condition of Claim 1 is satisfied by x^*. Let $\mathcal{Q}' \subseteq \mathcal{Q}(S)$. Then

$$\sum_{Q \in \mathcal{Q}'} r_Q = \sum_{Q \in \mathcal{Q}'} (1 - x^*(Q \setminus B(S))) = |\mathcal{Q}'| - \sum_{Q \in \mathcal{Q}'} x^*(Q \setminus B(S)).$$

According to Claim 1, what we have to show is that this quantity is at most $x^*(\cup_{Q \in \mathcal{Q}'} Q \cap B(S))$, which is exactly the content of Claim 2:

Claim 2 $x^*(\cup_{Q \in \mathcal{Q}'} Q) \geq |\mathcal{Q}'|$.

Indeed, denote $k := |\mathcal{Q}'|$, and let $C_0, \ldots, C_k$ the vertex sets of components of $S \setminus \cup_{Q \in \mathcal{Q}'} e_S^Q$, where $s \in C_0, t \in C_k$. Since $\{C_0, \ldots, C_k\}$ is a partition of V, every edge is counted twice in the sum of their coboundaries:

$$x^*(\cup_{Q \in \mathcal{Q}'} Q) = \frac{1}{2} \sum_{j=0}^{k} x^*(\delta(C_i)) \geq \frac{1}{2}(x^*(\delta(C_0)) + \sum_{j=1}^{k-1} x^*(\delta(C_j)) + x^*(\delta(C_k))) \geq \frac{1}{2}(1+2(k-1)+1) = k,$$

exactly as claimed, finishing the proof of Claim 2 and of Lemma 7. □

Proof of Theorem 4 and Restated Algorithm

We are now ready to prove Theorem 4 putting together Lemmas 6, 7 and 3. As a bonus, the proof reveals that the inner loop of our algorithm can be replaced by an appropiate definition of the objective function for the $T_F \triangle \{s, t\}$-join that anticipates the reconnection cost. We finish this section by developing this.

Proof of Theorem 4: Let $\mathcal{J}_F$ be a random $T_F \triangle \{s, t\}$-join with $\mathbb{E}[\mathcal{J}_F] = y_F$ (see Section 3) that is "derandomized" by the algorithm. We prove:

$$\mathbb{E}[c(2D_{F+\mathcal{J}_F})] \leq \sum_{Q \in \mathcal{Q}(S)} (x^*(Q) - 1) c(e_S^Q)$$

On the event $\mathcal{J}_F = J_F$ the least possible cost for reconnecting $F + J_F$ with pairs of doubled edges is $c(2D_{F+J_F})$. Another possible way of reconnecting is to add for all $b \in J_F \cap B(S)$, the lonely edges of $\mathcal{Q}(S, b)$ doubled, except one of them, as explained in (10) and the lines above it. Hence, we have

$$c(2D_{F+J_F}) \leq \sum_{b \in J_F \cap B(S)} \left(- \max_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q) + \sum_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q) \right). \tag{12}$$

Now we prove that the right hand side of this equation can be upper-bounded by the right hand side of the assertion of the theorem, by randomizing the choice of the unique cut in $\mathcal{Q}(S,b)$ for which the lonely edge is not used for reconnection, and then applying Lemma 6:

Note that the basic parity correction part of $y_F(b)$ is equal to $x^*(b)/2$ for $b \in B(S)$, and we thus need to show that the (PCE) and (PCL) vectors occurring for tree S are 0 for edges in $B(S)$.

This is immediately true for (PCE): e_S^Q for $Q \in \mathcal{Q}(S)$ is by definition contained in exactly one cut of $\mathcal{Q}(S)$, while the edges in $B(S)$ are those contained in at least two such cuts. The fact that this also holds for (PCL) is exactly the conclusion of Lemma 3. Hence, the condition of Lemma 6 that $Pr(b \in \mathcal{J}_F) \leq x^*(b)/2$ for every $b \in B(S)$ is satisfied.

Lemma 7 makes sure that one can define a probability distribution so that the (conditional) probabilities give a valid choice (by $(D0)$ and $(D1)$) and satisfy $(D2)$. So Lemma 6 can indeed be applied, that is, $\Pr(e_S^Q$ is used for reconnection $) \leq \frac{x^*(Q)-1}{2}$ for every $Q \in \mathcal{Q}(S)$. When e_S^Q is used for reconnection, it is doubled, so its contribution to the costs is $2c(e_S^Q)$, and therefore the expectation of the described random reconnection cost is $\sum_{Q \in \mathcal{Q}(S)}(x^*(Q)-1)c(e_S^Q)$. □

We now show how the counting of the above proof can be used to remove the inner loop (over all $J_F \in \mathcal{J}_F$) from our algorithm. Since the expected reconnection cost $2D_{F+J_F}$ can be upper bounded using (12) and it is in fact this value that is upper bounded in Theorem 4, we can replace $2D_{F+J_F}$ in the algorithm and in the summary of the analysis in Theorem 5 as well: rather than letting $2D_{F+J_F}$ be a doubled minimum cost spanning tree, we let it denote the edge set of which the cost is given by the right hand side of (12).

Redefining the objective function for the computation of a minimum cost $T_F \triangle \{s,t\}$-join in the following way can then replace the inner loop (over the T_F-joins of $\mathcal{J}_F$): for any $e \in E \setminus B(S)$ let $c'(e) = c(e)$, and for $b \in B(S)$ let

$$c'(b) := c(b) - \max_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q) + \sum_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q).$$

Then the minimum c'-cost $T_F \triangle \{s,t\}$-join J_F^* anticipates the cost of the reconnection with $2D_{F+J_F}$ as well. Since our analysis bounded the c'-cost of a convex combination of $T_F \triangle \{s,t\}$-joins this bound also applies to $c'(J_F^*)$.

We summarize the revised algorithm using the specifications of this section in Algorithm 2.

5 Layered Trees through Matroid Partition

In this section we prove that a layered $\mathcal{S}$, $\mathbb{E}[\mathcal{S}] = x^*$ exists, and can be found in a combinatorial way and strongly polynomial time. Recall that this means for $Q \in \mathcal{Q}(S)$ $(S \in \mathcal{S})$ and $x^*(Q) \geq x^*(Q')$, that Q' is also lonely.

Note that an arbitrary convex combination is not necessarily layered (see Figure 4).

Let the different cut-sizes of narrow cuts be $2 - \zeta_1 > 2 - \zeta_1 - \zeta_2 > \ldots > 2 - \zeta_1 - \ldots - \zeta_k = 1$,

$$\text{and } \mathcal{Q}_i := \{Q \in \mathcal{Q} : x^*(Q) \leq 2 - \zeta_1 - \ldots - \zeta_i\},$$

$$\mathcal{B}_i := \{S \subseteq E : S \text{ is a spanning tree of } G,\ |S \cap Q| = 1 \text{ for all } Q \in \mathcal{Q}_i\ \},$$

that is, $\mathcal{B}_i$ is the set of spanning trees of $G = (V,E)$ for which every cut $Q \in \mathcal{Q}_i$ contains exactly one edge. Clearly, $\sum_{i=1}^k \zeta_i = 1$, and $\mathcal{B}_1 \subseteq \ldots \subseteq \mathcal{B}_k$.

The following is a fractional (so weaker) version of Edmonds' matroid partition theorem in the form that we need it (see close versions to which it can be reduced [19, Chapter 42]):

Let x^* be an optimal solution to the subtour elimination LP for the $s-t$ path TSP.
Write x^* as a layered convex combination of spanning trees.
Let $\mathcal{S}$ be the collection of trees with positive coefficient.
for $S \in \mathcal{S}$ **do**
  Construct forest-based $\{s,t\}$-tour:
  $F := F(S) := S \setminus L(S)$.
  Let $c'(e) = c(e)$ for all $e \in E \setminus B(S)$.
  Let $c'(b) = c(b) + \sum_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q) - \max_{Q \in \mathcal{Q}(S,b)} 2c(e_S^Q)$. for all $b \in B(S)$.
  Let J_F^* be a minimum cost $T_F \triangle \{s,t\}$-join with respect to costs c'.
  Contract the components of $(V, F + J_F^*)$, and let $2D_{F+J_F^*}$ be the edge set of a doubled minimum cost spanning tree on the contracted graph.
  $P_1(S) = F + J_F^* + 2D_{F+J_F^*}$.
  Construct Christofides-based $\{s,t\}$-tour:
  $P_2(S) = S + J_S$, where J_S is a minimum cost $T_S \triangle \{s,t\}$-join with respect to costs c.
end for
Return a minimum cost $\{s,t\}$-tour from among $\bigcup_{S \in \mathcal{S}} \{P_1(S), P_2(S)\}$.

Algorithm 2: Restated Best-of-Many With Deletion

Theorem 8 *Let $M_i = (E, r_i)$ $(i = 1, \ldots, k)$ be matroids, $w \in \mathbb{R}^E$, and $\lambda_i \in \mathbb{R}$, $0 \le \lambda_i \le 1$ $(i = 1, \ldots, k)$. Let P^i be the convex hull of independent sets of M_i. There exist $x^i \in P^i$ for $i = 1, \ldots, k$ such that:*

$$\sum_{i=1}^{k} \lambda_i x^i = w, \tag{13}$$

if and only if for all $X \subseteq E$: $$\sum_{i=1}^{k} \lambda_i r_i(X) \ge w(X). \tag{14}$$

We say that edge $e \in E$ is in *layer* i ($i \in \{1, \ldots, k\}$), if there is a unique cut $Q \in \mathcal{Q}_i$ such that $e \in Q$. Edges may also be in several or no layers. Since $\mathcal{Q}_1 \supseteq \ldots \supseteq \mathcal{Q}_k$, the layers in which an edge is contained are consecutive numbers. (If it is contained exactly in layers $i \in [a,b] \cap \mathbb{N}$, then it is contained in at least 2 cuts of $\mathcal{Q}_i$ for $i < a$, and in none of the cuts $\mathcal{Q}_i$ for $i > b$.) Denote $L_i := \{e \in E : e \text{ is in layer } i\}$. So we can think of layer i to be the set L_i of "Gao edges" for the cuts $\mathcal{Q}_i$. The edges Gao used for his solution [13] were those of L_1.

Following and generalizing Gao [13] we call sets of the form $W \setminus U \neq \emptyset$ where $s \in U \subseteq W$, $Q = \delta(U) \in \mathcal{Q}$, $R = \delta(W) \in \mathcal{Q}$, *level-sets.* If in addition, $Q, R \in \mathcal{Q}_i$, $W \setminus U$ will be called a *level-set of layer i.* Gao [13, Lemma 3.2] proved that *level-sets induce connected subgraphs of G.* The natural ordering of chains induces a natural ordering of (inclusionwise) minimal level-sets starting with $\{s\}$ and ending with $\{t\}$. What Gao calls level sets are the minimal level sets of layer 1. In the Appendix we extend his connectivity result by proving that level-sets of $\mathcal{Q}_i$ are in fact $\zeta_1 + \ldots + \zeta_i$-connected (Lemma 15), and show the connection with layered convex combinations of spanning trees. We use here only that

$$L_i \cap Q \neq \emptyset \text{ for all } Q \in \mathcal{Q}_i \ (i = 1, \ldots, k) \text{ and the level-sets induce connected graphs.} \tag{15}$$

that can be checked anew, or as a special case of [13, Lemma 3.2], or apply Lemma 15 in Appendix D.

Denote $G_i := E \setminus \bigcup_{Q \in \mathcal{Q}_i} Q$. Let $M_i := (E, r_i)$ $(i = 1, \ldots, k)$ be the matroid whose independent sets are of the form $F_i \cup C_i$, where F_i is a forest in G_i and $C_i \subseteq L_i$, $|C_i \cap Q| \le 1$ for each $Q \in \mathcal{Q}^i$.

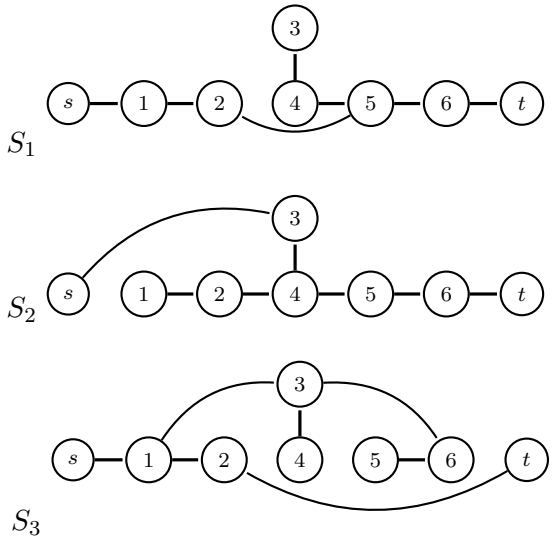

Figure 4: x^* from Figure 1 expressed as a convex combination of three spanning trees S_1, S_2, S_3, each with multiplier $\lambda_{S_i} = 1/3$, that is *not* layered. Note that $x^*(Q_j) = 5/3$ for $j = 2, \ldots, 5$, that there is exactly one S_i for which $|Q_j \cap S_i| = 1$, so in order for (9) to hold, Q_j has to be lonely in the unique tree S_i such that $|Q_j \cap S_i| = 1$. But then, for example, Q_2 is lonely in S_1, but Q_3 is not.

(Such sets form the independent sets of a matroid, as straightforward from the fact that the forests of any graph form the independent sets of a matroid.) The set of bases of M_i will be denoted by $\mathcal{B}_i$. Clearly, $B \in \mathcal{B}_i$ if and only if B meets each component of G_i in a spanning tree, and meets each cut of $\mathcal{Q}_i$ in exactly 1 edge. According to (15), $\mathcal{B}_i \neq \emptyset$. Summarizing:

$$M_i := (E, r_i) \text{ is a matroid, and its rank is } r_i(E) = n - 1 \ (i = 1, \ldots, k). \tag{16}$$

We first establish the connection between layered convex combinations and these matroids:

Lemma 9 *The following statements are equivalent:*

(i) *There exists a layered convex combination.*

(ii) *There exists a convex combination* $\sum_{S \in \mathcal{S}} \lambda_S S = x^*$ $(\lambda_S > 0$ *for* $S \in \mathcal{S}$, $\sum_{S \in \mathcal{S}} \lambda_S = 1)$ *of spanning trees, and* k *nonnegative integers* $1 \leq m_1 \leq \ldots \leq m_k = |\mathcal{S}|$, *so that the trees in the first* m_i *terms of the sum are in* $\mathcal{B}_i$, *and the sum of the first* m_i *coefficients is* $\sum_{j=1}^{i} \zeta_j$ *(for all* $i = 1, \ldots, k$*).*

(iii) *The equation (13) holds, where* $w = x^*$ *and* P^i *is the convex hull of* $\mathcal{B}_i$, $\lambda_i = \zeta_i$ $(i = 1, \ldots, k)$.

Furthermore, finding (i), (ii) or (iii) are (strongly) polynomially equivalent.

Proof: If (ii) holds, the sum of the coefficients from the first term and the m_1-th term is ζ_1, and the trees are in $\mathcal{B}_1$. From the $m_i + 1$-th until the m_{i+1}-th term the sum of the coefficients is $\sum_{j=1}^{i+1} \zeta_j - \sum_{j=1}^{i} \zeta_j = \zeta_{i+1}$, and the trees are in $\mathcal{B}_{i+1}$ $(i = 1, \ldots, k-1)$, by the constraint of (ii). This is exactly a combination of spanning trees in the form of (13), with the specifications of (iii). So (iii) is proved.

Consider now a convex combination in the form of (13) with the specifications of (iii). Substitute for x^i a convex combination of $\mathcal{B}_i$, and let $x^* = \sum_{i=1}^{k} \sum_{S \in \mathcal{B}_i} \lambda_S S$ be the corresponding convex combination of x^* into spanning trees, where $\zeta_i x^i = \sum_{S \in \mathcal{B}_i} \lambda_S S$.

$$\text{The set of lonely edges } L(S) \text{ for } S \in \mathcal{B}_i \text{ is defined as } S \cap L_i. \tag{17}$$

By the definition of L_i (5) is satisfied, moreover by (15) then $\mathcal{Q}(S) = \mathcal{Q}_i$. In order to prove (i), it remains to check that S is layered, and according to the definition end of Section 2.3 this means that $Q \in \mathcal{Q}_i$ and $x^*(Q') \leq x^*(Q)$ for a cut imply $Q' \in \mathcal{Q}_i$, which is obviously true.

Finally, suppose that (i) holds, and let $\sum_{S \in \mathcal{S}} \lambda_S S = x^*$ be a layered convex combination. In this sum, order the trees S in decreasing order of the largest cut-size among the cuts in $\mathcal{Q}(S)$. Then group the ties, that is, form k groups of sums (some of which may be empty) of trees whose lonely cuts have the same maximum cut size $2-\zeta_1, \ldots, 2-\zeta_1-\ldots-\zeta_i, \ldots, 2-\zeta_1, \ldots, 2-\zeta_1-\ldots-\zeta_k$ respectively. Since $\mathcal{S}$ is layered, for the trees S in the i-th group $\mathcal{Q}_i \subseteq \mathcal{Q}(S)$, and therefore in the first i groups there are only trees in $\mathcal{B}_i$.

Let m_i be the number of trees in the first i groups ($i = 1, \ldots, k$). By (9) which is satisfied by $L(S)$ (see the lines after (9) or the definition of layered at the end of Section 2.3) the probability that a cut of size $2-\zeta_1-\ldots-\zeta_i$ contains a lonely edge is exactly $\zeta_1+\ldots+\zeta_i$, and this happens exactly for the trees of the first i groups, so $m_i = \sum_{j=1}^{i} \zeta_i$, which provides a convex combination of the form of (ii). □

By Theorem 8 and Lemma 9, a layered convex combination exists if and only if (14) holds for our matroids, and then it can be determined in polynomial time:

Theorem 10 *The matroids M_i ($i = 1, \ldots k$) satisfy (14) with $w = x^*$, $\lambda_i = \zeta_i$, and r_i the rank function of the matroid $(E, \mathcal{B}_i)$; a layered convex combination exists and can be determined in strongly polynomial time, by running the matroid partition algorithm once.*

Proof: The existence of a convex combination satisfying Lemma 9 (ii) has already been proved using an interesting, but technically demanding local procedure, "tree-reassambling", by Gottschalk and Vygen [16], without the matroid property. So by Lemma 9 we have (iii) as well, and therefore, by the easy part of the matroid partition theorem, Theorem 8 this implies (14). Running Edmonds' matroid partition algorithm (see the references below) it cannot terminate with a violated condition, so it will terminate by finding a convex combination that satisfies (ii), which, again by Lemma 9 implies (i), that is, a layered convex combination, as claimed. □

Note that our proof uses the result of [16] to check that (14) holds; however, checking that (14) holds for a particular set of matroids is usually a straightforward task, that would provide a direct simpler proof for the existence of layered decompositions as well. Lemma 15 (Appendix D) shows a direct algorithmic proof for the existence of a basis-packing in M_i for all $i \in \{1, \ldots, k\}$ respecting the capacity function x^* with sum of coefficients $\zeta_1+\ldots+\zeta_i$. This is weaker than a layered convex combination: it is explained in Appendix D that having only one packing satisfying each of these k properties simultaneously would be sufficient for (and is equivalent to) a layered convex combination, that is, to (14) or (20). However, when this manuscript is submitted, we have not completed a direct approach for proving this.

For actually finding a layered convex combination through tree packing, the proofs are algorithmic provided a testing membership problem in tree polyhedra is available. ([19, Chapter 51, mainly 51.4]; or it may be better to use general matroid bases packing, see also Edmonds' "matroid base packing theorem"[9] (a straightforward consequence of Edmonds' matroid partition theorem) [12, 13.3.13], [19, Corollary 42.1d], and execute any weighted algorithm giving priority to smaller index matroids; for the polynomial solvability of matroid basis packing we refer to [7], and at [19, page 891] the exact competitive complexities can be checked.) The rank oracle can be computed in linear time for all our matroids $M_i = (E, r_i)$ as well: $r_i(X) = |V| - c(X \setminus \bigcup_{Q \in \mathcal{Q}_i} Q) + |Q \in \mathcal{Q}_i : Q \cap L_i \neq \emptyset|$, where $c(Y)$ ($Y \subseteq E$) denotes the number of connected components of (V, Y).

We conclude by summarizing the overall proof of our bounds from a bird's eye view: Theorem 10 allows us to avoid additional reconnection costs, by ensuring that the parity completion part of the

parity correction vector does not contribute to the probability of requiring reconnection across a lonely cut (Lemma 3). This is an important ingredient in the proof of Theorem 4, which bounds the total expected reconnection cost. Another important component of the bound of Theorem 4 is provided by Section 4, where we define a probability distribution that determines which lonely edges to use for reconnection, thus allowing us to "allocate" the cost of reconnection to the lonely cuts. Then Theorem 4 itself is used in the improved ratio that is our main result, Theorem 5.

References

[1] H. Ch. An, R. D. Kleinberg, and D. B. Shmoys. Improving Christofides' algorithm for the s-t path TSP. In: Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC'12). Revised in [2].

[2] H. Ch. An, R. D. Kleinberg, and D. B. Shmoys. Improving Christofides' algorithm for the s-t path TSP. *J. ACM*, 62(5):34, 2015.

[3] F. Barahona. Packing spanning trees. *Mathematics of Operations Research*, 20(1):104–115, 1995.

[4] J. Cheriyan, Z. Friggstad, and Z. Gao. Approximating minimum-cost connected t-joins. *Algorithmica*, 72(1):126–147, 2015.

[5] N. Christofides. Worst case analysis of a new heuristic for the traveling salesman problem. Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA, 1976.

[6] W. J. Cook. *In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation.* Princeton University Press, December 2011.

[7] W. Cunningham. Testing membership in matroid polyhedra, *Journal of Combinatorial Theory, Series B*, 36:161–188, 1984.

[8] G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2:393–410, 1954.

[9] J. Edmonds. Lehman's switching game and a theorem of Tutte and Nash-Williams, *J. of Research of the National Bureau of Standards,* Section B 69 (1965), pp. 125–130.

[10] J. Edmonds. Submodular functions, matroids and certain polyhedra. *In: Combinatorial Structures and Their Applications*, Proceedings of the Calgary International Conference on Combinatorial Structures and Their Applications 1969 (R. Guy, H. Hanani, N. Sauer, J. Schönheim, eds.), Gordon and Breach, New York 1970, pp. 69–87.

[11] J. Edmonds, and E. L. Johnson. Matching, Euler tours and the Chinese postman. *Mathematical Programming*, 5 (1973), 88–124.

[12] A. Frank. Connections in Combinatorial Optimization. Oxford University Press 2011.

[13] Z. Gao. An LP-based 3/2-approximation algorithm for the *s-t* path graph Traveling Salesman Problem. *Oper. Res. Lett.*, 41(6):615–617, 2013.

[14] Z. Gao. On the metric s-t path traveling salesman problem. *SIAM J. Discrete Math.*, 29(3):1133–1149, 2015.

[15] M. R. Garey, D. S. Johnson, and R. E.Tarjan. The planar hamiltonian circuit problem is np-complete. *SIAM Journal on Computing*, 5(4):704–714, 1976.

[16] C. Gottschalk and J. Vygen. Better *s-t* -tours by Gao trees. *CoRR*, abs/1511.05514, 2015.

[17] J. A. Hoogeveen. Analysis of Christofides' heuristic: Some paths are more difficult than cycles. *Operations Research Letters*, 10(5):291 – 295, 1991.

[18] C.St.J.A. Nash-Williams. Edge-disjoint spanning trees of finite graphs, *The Journal of the London Mathematical Society*, 36, 445–450, 1961.

[19] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

[20] A. Sebő. Eight-fifth approximation for the path TSP. In *Proceedings of the 16th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 362–374, 2013.

[21] A. Sebő, J. Vygen, Shorter tours by nicer ears: 7/5-approximation for graphic TSP, 3/2 for the path version, and 4/3 for two-edge-connected subgraphs, *Combinatorica*, **34** (5) October (2014), 597–629.

[22] A. Sebő. Problems about uniform covers, with tours and detours. In *Matematisches Forschungsinstitut Oberwolfach Report No. 51/2014, DOI: 10.4171/OWR/2014/51, November 2014*, on joint work with Y. Benchetrit and M. Stehlík, pages 2912–2915, March 2015.

[23] W.T. Tutte. On the problem of decomposing a graph into n connected factors, *The Journal of the London Mathematical Society*, 36, 221-230, 1961.

[24] J. Vygen. Reassembling trees for the traveling salesman. *CoRR*, abs/1502.03715, 2015.

[25] L. A. Wolsey. Heuristic analysis, linear programming and branch and bound. *Mathematical Programming Study*, 13:121–134, 1980.

APPENDIX

The appendix contains further illustrations of our techniques and additional results. In Appendix A, we illustrate our setting and analysis on the best-of-many Christofides' algorithm (without deletion) with a very short proof of the guarantee of 8/5. In Appendix B, we show that the deletion of lonely edges gives tight results in certain cases. In each of these part of the difficulties disappear, which allows to focus better on those that remain. In Appendix C, we give a self-contained proof of Claim 1, and in Appendix D we provide a new matroid construction inspired by Section 5. We also show a partial direct proof for the layered convex combination and exhibit what is missing from a full matroidal proof.

A Our Setting and Analysis for BOMC

Let us say that a narrow cut Q is small if $1 \leq x^*(Q) \leq \frac{3}{2}$, and large, if $\frac{3}{2} \leq x^*(Q) < 2$. (If $x^*(Q) = \frac{3}{2}$ it is both small and large.) We begin by considering the cases when all narrow cuts are small or all are large.

Throughout Appendix A, we take $L(S)$ to be the set of edges of S that are the unique edge in some narrow cut, i.e., (5) will hold with equality, and we can apply (8).

A.1 All Narrow Cuts Are Small

If all narrow cuts are small, then the probability that $|Q \cap \mathcal{S}|$ is even is at most $x^*(Q) - 1 \leq \frac{1}{2}$ by (8). Also, by (8), the vector $\frac{x^*}{2} + \sum_{Q \in \mathcal{Q}: |Q \cap S| \text{ even}} \frac{x^Q}{2}$ is in the $T_S \triangle \{s,t\}$-join polyhedron. Hence, the average or expected parity correction can be bounded by $\frac{1}{2}x^* + \sum_{Q \in \mathcal{Q}} \frac{1}{2}\frac{x^Q}{2} \leq \frac{1}{2}x^* + \frac{1}{4}p^*$, where the inequality follows by noting that if e is the lonely edge in narrow cut Q in some tree S, then e must be on the $s-t$ path in S, so $\sum_{Q \in \mathcal{Q}} x^Q \leq \mathbb{E}[\mathcal{S}(s,t)] = p^*$. So, if all narrow cuts are small, we can again bound the cost of the best-of-many solution by $\frac{3}{2}c(x^*) + \frac{1}{4}c(p^*)$, and combined with Observation 2, this implies a bound of $\frac{8}{5}c(x^*)$ for this case as well.

A.2 All Narrow Cuts Are Large

If all narrow cuts are large, then the basic parity correction vector (6) with $\gamma = \frac{1}{4}$ yields already as good a vector as the best one in the "All Small" case: it is in the $T_S \triangle \{s,t\}$-join polyhedron for *any* S, since by assumption $x^*(Q) \geq \frac{3}{2}$ for all Q.

Hence, the average cost of the spanning tree plus parity correction is at most $c(x^*) + \frac{1}{2}c(x^*) + \frac{1}{4}c(p^*)$. Using Observation 2, we can also bound the cost by $2c(x^*) - c(p^*)$, and since $\min\{\frac{3}{2}c(x^*) + \frac{1}{4}c(p^*), 2c(x^*) - c(p^*)\} \leq \frac{8}{5}c(x^*)$, we get the desired result.

A.3 Combining the Two Cases

It turns out that the two extreme cases described above can be "merged" by essentially taking the average of the two, and surprisingly without having any extra cost. We show two, slightly different proofs of the ratio 8/5 of [20]: the first, applying systematically the setting of the present paper, the second, even simpler, pointing more at the intuitive "averaging merge" also constituting a bridge to the original solution, that we include as a remark.

Connectivity: Tree $\mathcal{S}$.

Expected Value of Connectivity: $\mathbb{E}[S] = x^*$.

Parity Correction: The cost of a $T_S \triangle \{s,t\}$-join can be upper bounded by the cost of a point in the $T_S \triangle \{s,t\}$-join polyhedron, which is the sum of the following two vectors:

Basic Parity Correction (BP): $\frac{1}{2}x^* + \gamma\, S(s,t)$. *Expected Value of (BP)*: $\frac{1}{2}x^* + \gamma\, p^*$.

Parity Completion for Even Cuts (PCL): $\left(\frac{2-2\gamma-x^*(Q)}{2-x^*(Q)}\right)\frac{x^Q}{2}$.

We add a parity completion term for narrow cut Q and tree $S \in \mathcal{S}$ only if $|Q \cap S|$ is even, and $x^*(Q) \in [1, 2(1-\gamma))$. By (8), (BP)+(PCL) is a vector in the $T_S \triangle \{s,t\}$-join polyhedron for each $S \in \mathcal{S}$. Since we add (PCL) for Q with probability $1 - x^Q(Q) \leq x^*(Q) - 1$ by (8), the expectation of the (PCL) vector is at most $\sum_{Q \in \mathcal{Q}: x^*(Q) < 2(1-\gamma)} \frac{(1-\gamma-x^*(Q)/2)\,(x^*(Q)-1)}{2-x^*(Q)} x^Q$.

For each $Q \in \mathcal{Q}$ the coefficient of x^Q in this sum is at most

$$\left(1-\gamma-\frac{x}{2}\right)\frac{x-1}{2-x},$$

where $x = x^*(Q)$. For fixed γ, the unique maximum of this expression in the interval $[1,2]$ is $x = 2-\sqrt{2\gamma}$. Substituting this in the formula we find that the unique minimum is $\gamma = \frac{1}{8}$, when the worst cut-size is $x = 2 - \sqrt{2 \times \frac{1}{8}} = \frac{3}{2}$, the unique intersection of small and large cut-sizes. Substituting these values into our bound, we get that the coefficient of x^Q for all $Q \in \mathcal{Q}$ is at most $\frac{1}{8}$.

Finally note that if e is the lonely edge in narrow cut Q in some tree S, then e must be on the $s-t$ path in S, so (PCL) is finally bounded by $\frac{1}{8}\sum_{Q \in \mathcal{Q}} x^Q \leq \frac{1}{8}\mathbb{E}\left[\mathcal{S}(s,t)\right] = \frac{1}{8}p^*$.

Expected Value of (PCL): $\leq \frac{1}{8}p^*$.

Expected Value of Parity Correction: $\frac{1}{2}x^* + \frac{1}{4}p^*$.

The total expected cost of the solution can thus again be bounded by $\frac{3}{2}c(x^*) + \frac{1}{4}c(p^*)$. Combined with Observation 2, we can bound the cost of the constructed solution by $\frac{8}{5}c(x^*)$.

Remark: We can also arrive at this solution without optimizing over γ, with less computation:

For the bounds in the "all small" and "all large" case the unique worst cut-size is $x = \frac{3}{2}$. Take a convex combination (for instance with $\frac{1}{2}$ coefficients, leading to their arithmetic mean), of the two parity correction vectors π_1 and π_2 that "even cuts" receive in the two cases respectively: $\pi_1 := \frac{1}{2}x^* + \sum \frac{x^Q}{2}$ and $\pi_2 := \frac{1}{2}x^* + \frac{1}{4}p^*$. Then $\pi := \frac{\pi_1+\pi_2}{2} = \frac{1}{2}x^* + \frac{1}{2}\sum \frac{x^Q}{2} + \frac{1}{8}p^*$ is certainly a valid parity completion vector for even cuts of size $\frac{3}{2}$, as it is the convex combination of two valid parity correction vectors.

In fact, since π contains the term $\frac{1}{8}p^*$, the vector is automatically valid for even cuts of size at least $\frac{7}{4}$. Moreover, if we replace π by $\frac{1}{2}x^* + \sum\left(\frac{7/4-x^*(Q)}{2-x^*(Q)}\right)\frac{x^Q}{2} + \frac{1}{8}p^*$ (where we have to add the middle term only for even cuts), then we get a valid parity correction vector. By (8), we add the middle term with probability $x^*(Q) - 1$ for each $Q \in \mathcal{Q}$, so the mean value of parity correction is at most

$$\frac{1}{2}x^* + \sum_{Q \in \mathcal{Q}} (x^*(Q)-1)\left(\frac{7/4-x^*(Q)}{2-x^*(Q)}\right)\frac{x^Q}{2} + \frac{1}{8}p^*$$

Now, for $1 \leq x < 2$,

$$\left(\frac{7}{8}-\frac{x}{2}\right)\frac{x-1}{2-x} \leq \frac{1}{8}.$$

(Multiply both sides of the inequality by $2-x > 0$, $(\frac{7}{8}-\frac{x}{2})(x-1) - \frac{1}{8}(2-x) = -\frac{x^2}{2} + \frac{3x}{2} - \frac{9}{8} = -\frac{1}{8}(2x-3)^2 \leq 0$.) So we got the same bound as in the two special cases.

B Tight Bound on the Integrality Gap for Special Cases

We show some simple cases where the ratio and integrality gap 3/2 can be proved.

B.1 Disjoint and Almost-Disjoint Narrow Cuts

Let $G = (V, E)$ be a graph and denote $\mathbf{1}$ be the all-1 vector on the edges. A naturally arising open question relaxing the integrality gap of the subtour elimination LP is to determine the smallest constant α so that $\alpha\mathbf{1}$ is *feasible*, that is, a convex combination of T-tours[2] [22]:

For 3-edge-connected graphs $\mathbf{1}$ is always feasible. Indeed, setting 2/3 on all the edges is in the (dominant of the) convex hull of spanning trees and 1/3 is in the (dominant of the) convex hull of T-joins for every T. It has been pointed out that this result is not best possible. Indeed, the conjecture that the integrality gap of the subtour elimination LP is $\frac{4}{3}$ for the metric TSP is well-known to be equivalent to the feasibility of $\frac{4}{3}x$ provided x is in the subtour polytope. Therefore, if this conjecture is true, $\alpha = \frac{4}{3} \cdot \frac{2}{3} = \frac{8}{9}$ is also feasible for $T = \emptyset$, which may be a weaker and more hopeful conjecture to prove than the $\frac{4}{3}$-conjecture about the integrality gap.

However, if $T \neq \emptyset$ the integrality gap is at least 3/2 (see [2] for an example for $T = \{s, t\}$), and it follows that $\alpha \geq 1$ is necessary for $\alpha\mathbf{1}$ to be feasible. Instead of 3-edge-connectivity we have to require then only that T-even cuts (that is, cuts $\delta(U)$, $U \subseteq V$, $|T \cap U|$ is even) are of size at least 3. The conjecture that the integrality gap of the subtour elimination LP is $\frac{3}{2}$ implies that $\alpha = \frac{3}{2} \cdot \frac{2}{3} = 1$ already defines a feasible vector, that is, $\mathbf{1}$ is feasible for such graphs.

Our techniques allow us to prove shortly that this is indeed true, that is:
If all T-even cuts of G are of size at least 3, $\mathbf{1}$ is a convex combination of T-tours. For T-tours, narrow cuts are also defined, and have similarly nice properties. They do not form a chain, but they do form a laminar system [4]. Since we will see that we do not need reconnection here, our proof can be extended to this case, but for ease of communication we restrict ourselves to $T = \{s, t\}$ in the following. Then the condition means just that identifying s and t we get a 3-edge-connected graph. The generalization for arbitrary T is straightforward.

Theorem 11 *Let $G = (V, E)$ be a connected graph, $s, t \in V$ and let $|\delta(U)| \geq 3$ for any $|U \cap \{s, t\}| \neq 1$. Then, $\mathbf{1}$ can be expressed as a convex combination of $\{s, t\}$-tours.*

Proof: We can suppose that G is 2-edge-connected, since a cut-edge necessarily separates s and t, and we can proceed by induction, separating the problem into two. Note also that $x^*(e) = \frac{2}{3}$ on every $e \in E$ is a feasible solution to the subtour elimination LP for $s - t$ path TSP. Hence, we can use the BOMD algorithm on x^*, where the narrow cuts are exactly the $s - t$ cuts of G containing two edges. We will show that the forest-based $\{s, t\}$-tours $F + J_F$ give us the desired convex combination; in other words, that reconnection is never needed.

Forest: We take an arbitrary decomposition of x^* into spanning trees $S \in \mathcal{S}$, and let $F := F(S) := S \setminus L(S)$.
Expected Value of Forest: $x^* - \sum_{Q \in \mathcal{Q}} x^Q$.

Parity Correction (PC): We take y_F to be the basic parity correction vector with $\gamma = 0$, and we add as parity completion $x^Q/2$ for every narrow cut Q such that $|Q \cap F|$ is even. As observed in Section 2.2.3, this indeed gives a feasible parity correction vector.
Expected Value of (PC): $\frac{1}{2}x^* + \frac{1}{2}\sum_{Q \in \mathcal{Q}} x^Q$.

[2] A T-tour is a connected graph in which the set of odd degree vertices is T.

Reconnection: Note that the fact that the non-$s-t$-cuts of G have at least three edges implies that *the narrow cuts are pairwise disjoint*:

Indeed, suppose there exist narrow cuts $Q_1 = \delta(U_1), Q_2 = \delta(U_2)$ (each containing exactly two edges), $s \in U_1$, $t \in U_2$ and $Q_1 \cap Q_2 \neq \emptyset$ Then the left hand side of (4) is $2 \cdot \frac{2}{3} + 2 \cdot \frac{2}{3}$, and the first term of the right hand side is $3 \cdot \frac{2}{3} = 2$ by the condition of the theorem; the last term on the right hand side is at least $2 \cdot \frac{2}{3}$ because a common edge, that exists by the non-emptiness asumption, has one endpoint in $U_1 \setminus U_2$ and the other in $U_2 \setminus U_1$. Since $\frac{4}{3} \neq \frac{5}{3}$ this is a contradiction.

It follows that there are no bad edges, since these are contained in at least two narrow cuts:

Expected Value of Reconnection: 0.

By adding up the different parts, we see that $\frac{3}{2}x^* - \frac{1}{2}\sum_{Q \in \mathcal{Q}} x^Q \leq \mathbf{1}$ is already in the convex hull of $\{s,t\}$-tours. □

We actually proved that $\frac{3}{2}x^* - \frac{1}{2}\sum_{Q \in \mathcal{Q}} x^Q$ *is a convex combination of* $\{s,t\}$ *tours, under the more general condition that the narrow cuts are disjoint.* We can strengthen this result as follows.

Theorem 12 *Let x^* be a feasible solution to the subtour elimination LP for $s-t$ path TSP, and let $G = (V,E)$ be the support graph of x^*. If no edge in G is contained in more than two narrow cuts, then $\frac{3}{2}x^*$ is a convex combination of $\{s,t\}$-tours.*

Proof: By Lemma 1, we know that there exist $s \subseteq U_1 \subset U_2 \subset U_2 \subset \ldots U_k$ such that the narrow cuts are exactly $\delta(U_i)$ for $i = 1, \ldots, k$. We will call a narrow cut $Q = \delta(U_i)$ odd-numbered or even-numbered, depending on whether i is odd or even. We decompose x^* into spanning trees. However, we will now only remove lonely edges from either the odd-numbered narrow cuts, or the even-numbered narrow cuts, each with probability $\frac{1}{2}$. This changes the expectation of the forest to $x^* - \frac{1}{2}\sum_{Q \in \mathcal{Q}} x^Q$.

The upper bound $\frac{1}{2}x^* + \frac{1}{2}\sum_{Q \in \mathcal{Q}} x^Q$ on (PC) is still valid (see above), and we are again guaranteed that $2D_{F+J_F} = \emptyset$, since a bad edge can be contained in only two consecutive narrow cuts, and we never delete the lonely edges of both of these. Whence the cost of reconnection is 0. Adding up the two nonzero contributions we get exactly $\frac{3}{2}x^*$. □

B.2 A 3/2-Approximation if All Narrow Cuts Are Small

In this section we analyze the BOMD algorithm when all narrow cuts Q have $x^*(Q) \leq \frac{3}{2}$. The following inequality that appears in Vygen [24] is very handy to exploit such upper bounds, and gives an easier (but weaker) tool than the tools developed in Section 4 for bounding reconnection costs:

$$\text{For any } Q_1, Q_2 \in \mathcal{Q},\ x^*(Q_1 \cap Q_2) \leq \frac{1}{2}(x^*(Q_1) + x^*(Q_2)) - 1. \tag{18}$$

We derive this by stopping before the last inequality in the proof of Claim 2, with the particular choice $\mathcal{Q}' = \{Q_1, Q_2\}$. Indeed, then $x^*(\cup_{Q \in \mathcal{Q}'} Q) \geq |\mathcal{Q}'| - 1 + \frac{1}{2}(x^*(Q_1) + x^*(Q_2)) = 1 + \frac{1}{2}(x^*(Q_1) + x^*(Q_2))$. Subtracting this inequality from $x^*(Q_1) + x^*(Q_2) = x^*(Q_1) + x^*(Q_2)$, we get (18).

In our special case $x^*(Q_1 \cap Q_2) \leq \frac{1}{2}$ follows for all $Q_1, Q_2 \in \mathcal{Q}$.

We show below that this inequality is sufficient to get a good bound on the cost of reconnection, assuming that we use a layered convex combination from Section 5. In fact, we get the following simple analysis of the forest-based solutions found by the BOMD algorithm:

Forest: Forest $F := F(S) := S \setminus L(S)$.

Expected Value of Forest: $x^* - \sum_{Q \in \mathcal{Q}} x^Q$.

Parity Correction: The fractional solution y_F in the $T_F \triangle \{s,t\}$-join polyhedron is constructed by adding the following vectors:

Basic Parity Correction (BP): $\frac{1}{2}x^*$.
Note that $\frac{1}{2}x^*(Q) \geq 1$ for any cut Q that is *not* narrow, so (BP) is enough to ensure that y_F satisfies the constraints of the $T_F \triangle \{s,t\}$-join polyhedron for the cuts that are not narrow.
Expected Value of (BP): $\frac{1}{2}x^*$.

Parity Completion for Empty Cuts (PCE): $\left(1 - \frac{x^*(Q)}{2}\right) e_S^Q$.
Note that (BP)+(PCE) are indeed sufficient to ensure that y_F satisfies the constraints of the $T_F \triangle \{s,t\}$-join polyhedron for narrow cuts Q such that $|Q \cap F| = 0$.
Expected Value of (PCE): $\sum_{Q \in \mathcal{Q}} \left(1 - \frac{x^*(Q)}{2}\right) x^Q$.

Parity Completion for Non-Empty Even Cuts (PCL): $\frac{1}{2}x^Q$.
Recall from Section 2.2.3 that indeed (BP) + (PCL) are indeed sufficient to ensure that y_F satisfies the constraints of the $T_F \triangle \{s,t\}$-join polyhedron for narrow cuts Q such that $|Q \cap F|$ is even and ≥ 2. We also have that $\Pr(|Q \cap \mathcal{F}| \geq 2) \leq \Pr(|Q \cap \mathcal{S}| \geq 2) \leq x^*(Q) - 1$, and thus:
Expected Value of (PCL): $\frac{1}{2} \sum_{Q \in \mathcal{Q}} (x^*(Q) - 1)\, x^Q$.

Expected Value of Parity Correction: $\frac{1}{2}x^* + \frac{1}{2} \sum_{Q \in \mathcal{Q}} x^Q$.

Reconnection: When all cuts are small, we can use a simple rule to implement reconnection (10): for each bad edge b in $\mathcal{J}_F$, we choose the "leftmost" of the lonely cuts Q in $\mathcal{Q}(S,b)$ as the cut for which we do *not* add a doubled edge. (Cut $Q' = \delta(U') \in \mathcal{Q}$ is to the left of $Q = \delta(U) \in \mathcal{Q}$, if $s \in U' \subset U$; by Lemma 1 this is a total order on the narrow cuts.)

Now let $Q \in \mathcal{Q}(S)$, and compute the probability of adding two copies of e_S^Q for reconnection. For this to happen, $\mathcal{J}_F$ must contain an edge $b \in Q' \cap Q$ where Q' is the first lonely cut in S to the left of Q. The probability that $\mathcal{J}_F$ contains such an edge is the sum of basic parity correction on $Q' \cap Q$, plus parity completion on $Q' \cap Q$: the former of these two is at most $\frac{x^*(Q' \cap Q)}{2} \leq \frac{1}{4}$ by (18) under our assumption that all cuts are small; the latter is 0 by Lemma 3, since $\mathcal{S}$ is layered and the parity completion vectors are scalar multiples of x^Q.

Hence, for a given S, we get a bound on the reconnection cost of $\frac{1}{2} \sum_{Q \in \mathcal{Q}(S)} c(e_S^Q)$, that is:

Expected Value of Reconnection Cost: $\mathbb{E}\left[c(2D_{\mathcal{F}+\mathcal{J}_\mathcal{F}})\right] \leq \sum_{Q \in \mathcal{Q}} \frac{1}{2} c(x^Q)$.

We thus have that the total mean cost of $P_1(\mathcal{S})$ can be bounded by

$$\left(x^* - \sum_{Q \in \mathcal{Q}} x^Q\right) + \left(\frac{1}{2}x^* + \sum_{Q \in \mathcal{Q}} \frac{1}{2}x^Q\right) + \sum_{Q \in \mathcal{Q}} \frac{1}{2}x^Q = \frac{3}{2}x^*.$$

One not-small narrow cut size

We can extend the above analysis to the case when there exists one value $z > \frac{3}{2}$ such that for all narrow cuts Q, either $x^*(Q) \leq \frac{3}{2}$ or $x^*(Q) = z$. In this case, we let the parameter γ in the basic parity correction vector be non-zero, where the exact value is determined by the analysis. Compared to the ratio 3/2 we lose 1/34, like in the general case!

Now we use the same rule to deal with the tree S in the event $b \in B(S) \cap \mathcal{J}_F$: choose the leftmost lonely cut in $\mathcal{Q}(S,b)$ as the cut for which we do not use the edge e_S^Q for reconnection. As in the case when all cuts are small (see the above proof), let $Q \in \mathcal{Q}(S)$ and consider the probability of using e_S^Q for reconnection. Again, for this to happen, $\mathcal{J}_F$ must contain an edge $b \in Q' \cap Q$ where

Q' is the first lonely cut in S to the left of Q. The probability that $\mathcal{J}_F$ contains such an edge is the sum of basic parity correction on $Q' \cap Q$, plus parity completion on $Q' \cap Q$: the former of these two is now at most $\frac{x^*(Q' \cap Q)}{2} \leq \frac{1}{2}(x^*(Q) - 1)$ by (18); the latter is 0 by Lemma 3, since $\mathcal{S}$ is layered and the parity completion vectors are scalar multiples of x^Q.

This is exactly the bound on the probability of adding two copies of e_S^Q shown for the general case in Section 4, but the proof is much simpler and combinatorial, it is ignoring Section 4!

For a given tree S, and its forest $F = S \backslash L(S)$, we can thus bound $\mathbb{E}\left[c(2D_{F+\mathcal{J}_F})\right] \leq \sum_{Q \in \mathcal{Q}(S)}(x^*(Q)-1)c(e_S^Q)$, and taking the average of all $S \in \mathcal{S}$, we get a bound of $\sum_{Q \in \mathcal{Q}}(x^*(Q) - 1)x^Q$.

The remainder of the computation is exactly the same as in Section 3.

We are not able to improve the obtained bound, even not if $z = \frac{7}{4}$ and the only other cut-size is 1 and occurs only for $\delta(s)$ and $\delta(t)$. Being able to reach in general the best bound we know for this easy case we reach the maximum within the possibilities of our methods.

C A Self-Contained Proof of the Extended Kőnig-Hall Theorem

In this section we give a self-contained proof of Claim 1 (see proof of Lemma 7).

We use the duality theorem of linear programming (instead of Farkas' Lemma), so we add the objective function 0 to minimize. The dual of the LP defined by this objective function constrained by the inequalities $(D0)$, $(D1)$, $(D2)$, is simple to state, since every variable of the primal linear program is contained with nonzero coefficients, in at most one inequality of $(D1)$ and one inequality of $(D2)$:

Denoting by y_Q $(Q \in \mathcal{Q}(S))$, y_b, $(b \in B(S))$ the dual variables associated to the inequalities in $(D1)$, $(D2)$ respectively, after multiplying $(D1)$ by -1, the dual is:

$$\text{Maximize } \sum_{Q \in \mathcal{Q}(S)} r_Q y_Q - \sum_{b \in B(S)} y_b, \text{ subject to:}$$

$$x^*(b)y_Q - y_b \leq 0, \text{ and } y_Q \geq 0,\ y_b \geq 0 \text{ for all } b \in B(S), Q \in \mathcal{Q}(S,b).$$

Clearly, the dual LP is feasible, and by the duality theorem of linear programming, $(D0)$, $(D1)$, $(D2)$ is feasible if and only if the optimum of the dual LP is zero; furthermore, note that for an arbitrary choice $y_Q \geq 0 (Q \in \mathcal{Q}(S))$, the optimal choice y_b $(b \in B(S))$ is uniquely determined as

$$y_b := x^*(b) \max_{Q \in \mathcal{Q}(S,b)} y_Q$$

In other words $(D0)$, $(D1)$, $(D2)$ is feasible, if and only if

$$\text{(Farkas)} \qquad \sum_{Q \in \mathcal{Q}(S)} r_Q y_Q \leq \sum_{b \in B(S)} x^*(b) \max_{Q \in \mathcal{Q}(S,b)} y_Q, \text{ for all } y_Q \geq 0 (Q \in \mathcal{Q}').$$

Note that the condition of Claim 1 is exactly the $y_Q := 1$ $(Q \in \mathcal{Q}')$ special case of (Farkas)! So to prove Claim 1, we only have to check that whenever (Farkas) is violated the condition Claim 1 is also violated.

To prove this, suppose by contradiction that there exists $y_Q (Q \in \mathcal{Q}(S))$ be such that (Farkas) is violated, that is,

$$\text{(violatedFarkas)} \qquad \sum_{Q \in \mathcal{Q}(S)} r_Q y_Q > \sum_{b \in B(S)} x^*(b) \max_{Q \in \mathcal{Q}(S,b)} y_Q,$$

We suppose the number of different values of y_Q is minimum here, and let $\eta =: \max\{y_Q : Q \in \mathcal{Q}\}$. Note that if all $y_Q > 0$ are equal (to η) then dividing all y_Q ($Q \in \mathcal{Q}(S)$) by η, we see that Claim 1 is violated also.

If they are not all equal, let the second largest value of y_Q be $\eta - \varepsilon$ with $\eta > \varepsilon > 0$ and let $\mathcal{Q}'' := \{Q \in \mathcal{Q}(S) : y_Q = \eta\} \subsetneq \mathcal{Q}'$,. Now if the condition in Claim 1 is not violated then

$$\sum_{Q \in \mathcal{Q}''} r_Q \leq x^*(\cup_{Q \in \mathcal{Q}''} Q \cap B(S)).$$

Subtracting ε times the left hand side of this inequality from (violated Farkas), it is straightforward to see that (as usual, with $y_b := x^*(b) \max_{Q \in \mathcal{Q}(S,b)} \varepsilon$ the right hand side of (violated Farkas) is also decreased by at least ε times the right hand side. We get that (Farkas) is violated with one fewer y_Q value, a contradiction, finishing the proof of the claim.

D Matroids

We discuss here a slightly different way of handling layered convex combinations by using basis packing rather than the matroid partition approach of Section 5. These two are well-known to be equivalent, but both may have their advantage at use. An advantage of basis packing for us is that it helps presenting a partial direct proof of Theorem 10, and measure what is missing for a full proof in a clean way; see Section D.1). In the second subsection, we show a new matroid construction generalizing the matroids of Section 5, which may provide some additional insight.

D.1 How to Check for Basis Packing, and What Is Missing?

We state here a version of the "Basis Packing Theorem" of Edmonds for graphs. This provides some insight to the graphs induced by the level sets of the layers, implying a somewhat weaker statement than the layered convex combination of spanning trees, and making clear what is missing from a full direct proof (as promised in Section 5).

For a given vector w to be a convex combination of spanning trees is the $\zeta = 1$ special case of the following well known result of Tutte [23] and Nash-Williams [18] about packing spanning trees in graphs, itself a special case of matroid partition. This is a fractional (so weaker) version [19, Theorem 51.1a]:

For family $\mathcal{P}$ of disjoint subsets of V, let $\delta(\mathcal{P})$ denote the set of edges in E whose two endpoints are in two different sets of $\mathcal{P}$.

Theorem 13 *Let $G = (V, E)$ be a graph, $\mathcal{B}$ the set of its spanning trees, $w \in \mathbb{R}^E$, and $\zeta \in \mathbb{R}_+$. There exist $\lambda_B \in \mathbb{R}_+$ ($B \in \mathcal{B}$, $|\{B \in \mathcal{B} : \lambda_B > 0\}| \leq |E|$) such that:*

$$\sum_{B \in \mathcal{B}} \lambda_B B \leq w, \ \sum_{B \in \mathcal{B}} \lambda_B = \zeta, \tag{19}$$

if and only if for any partition $\mathcal{P}$ of V: $w(\delta(\mathcal{P})) \geq \zeta(|\mathcal{P}| - 1)$, (20)

and λ_B ($B \in \mathcal{B}$) or a violated condition can be found in strongly polynomial time.

We prove that for all $i = 1, \ldots, k$ (19) holds for $\mathcal{B} := \mathcal{B}_i$ and $\zeta = \zeta_1 + \ldots + \zeta_i$, and we can actually go to the furthest possible statement in this direction. For each basis in the chain $\mathcal{B}_1 \subseteq \ldots \subseteq \mathcal{B}_k$ this graph version of the basis packing theorem is exactly what we need.

Still, we also state basis packing theorem for matroids whose condition would be equivalent to check to that of Theorem 8:

Theorem 14 *Let $M_i = (E, r_i)$ $(i = 1, \ldots, k)$ be matroids, $w \in \mathbb{R}^E$, $R := r_1(E) = \ldots = r_k(E) = x^*(E)$, and $\lambda_i \in \mathbb{R}$, $0 \leq \lambda_i \leq 1$ $(i = 1, \ldots, k)$. Let P^i the convex hull of $\mathcal{B}_i$ the set of bases of M_i. There exist $x^i \in P^i$ for $i = 1, \ldots, k$ such that:*

$$\sum_{i=1}^{k} \lambda_i x^i = w, \tag{21}$$

$$\text{if and only if for all } X \subseteq E : w(X) \geq \textstyle\sum_{i=1}^{k} \lambda_i (R - r_i(X)). \tag{22}$$

and the coefficients of a convex combination or a violated condition can be found in strongly polynomial time.

This is the fractional version of [9] [19, Corollary 42.1d], a generalization of Theorem 13 for arbitrary matroids. A natural condition to check would be the condition of this theorem for our matroids. However, this appears to be a somewhat confusing task that we could not fully solve so far.

Lemma 15 *For each $i = 1, \ldots k$, $\mathcal{B}_i$ satisfies (19) with $w := x^*$, and $\zeta := \zeta_1 + \ldots + \zeta_i$.*

Proof: It is necessary and sufficient to check two inequalities: (20) for the graphs induced by level-sets of $\mathcal{Q}_i$ and $x^*(Q) \geq \zeta_1 + \ldots + \zeta_i$ for cuts $Q \in \mathcal{Q}_i$. Then we get by Theorem 13 a "packing" of ζ spanning trees in the level sets, and the same quantity of edges in the cuts of $\mathcal{Q}_i$; these can be straightforwardly pasted together to get the claimed packing. Both inequalities are consequences of the following claim:

Claim: Let $\mathcal{P}$ be a partition of $W \setminus U \neq \emptyset$ where $s \in U \subseteq W$, $Q = \delta(U)$, $R = \delta(W)$, $x^*(Q), x^*(R) \leq 2 - \zeta$. Then $x^*_{W \setminus U}(\delta(\mathcal{P})) \geq \zeta(|\mathcal{P}| - 1)$, where $x^*_{W \setminus U}$ is a restricton of x^* to the set of edges induced by $W \setminus U$.

Indeed, applying the submodular inequality (4) to W and $V \setminus U$:

$$2(2 - \zeta) \geq x^*(Q) + x^*(\delta(R) \geq x^*(\delta(W \setminus U)), \tag{23}$$

where $x^*(\delta(W \setminus U)) = \sum_{P \in \mathcal{P}} x^*(\delta(P)) - 2x^*_{W \setminus U}(\delta(\mathcal{P})) \geq 2|\mathcal{P}| - 2x^*_{W \setminus U}(\delta(\mathcal{P}))$. Substituting this into (23) we obtain:

$$x^*_{W \setminus U}(\delta(\mathcal{P})) \geq |\mathcal{P}| - 2 + \zeta. \tag{24}$$

Now if $|\mathcal{P}| = 2$ this is exactly the claimed inequality, and in case $|\mathcal{P}| > 2$ it is stronger. The claim is proved.

Now the Lemma also follows by noting that in case of $Q, R \in \mathcal{Q}_i$, the condition of the Claim is satisfied with $\zeta := \zeta_1 + \ldots + \zeta_i$. Moreover, if U and W are in addition such that there exists exactly one Y such that $U \subsetneq Y \subsetneq W$, $\delta(Y) \in \mathcal{Q}_i$, then apply the claim to the partition of size 2 $\{Y \setminus U\}$, to get $x^*(L_i \cap \delta(Y)) \geq \zeta$, with the same ζ.

Pasting together the ζ spanning trees of inclusionwise minimal level sets of $\mathcal{Q}_i$ and of the ζ edges in $L_i \cap Q$ for each $Q \in \mathcal{Q}_i$ provides the packing of ζ bases. □

Note that $|\mathcal{P}| - 2 + \zeta = \zeta(|\mathcal{P}| - 1) + (1 - \zeta)(|\mathcal{P}| - 2)$. With this remark (24) and Theorem 14 straightforwardly imply the following stronger result:

Theorem 16 *For every $i = 1, \ldots k$ there exists a convex combination of forests majorated by x^* consisting of a packing of $\zeta = \zeta_1 + \ldots + \zeta_i$ spanning trees that are bases of M_i and $\zeta_{i+1} + \ldots + \zeta_k$ forests, each with at most two components induced in each level set of M_i.*

Such a convex combination is the best possible in terms of spanning trees and graphs with at most two components in each level set. However, we do not know how to use this sharpening to have all the packings simultaneously.

Recall $\mathcal{B}_1 \subseteq \ldots \subseteq \mathcal{B}_k$. A convex combination of spanning trees containing simultaneously all the packings of Theorem 16 is exactly the assertion of Lemma 9 (ii), and as such, by Lemma 9 is equivalent to the existence of a layered convex combination. Unfortunately, Theorem 16 is weaker than (iii).

D.2 A (New?) Matroid Construction

We present the matroid construction of Section 5 starting from an arbitrary matroid. Is this matroid construction known? The graph special case was very useful for us, and the matroidal proofs may provide some more insight.

A *cut* in a matroid is a minimal set that does not intersect any circuit in exactly 1 element; a set of elements is a cut if and only if it is a circuit of the dual matroid. A matroid will be given in the form $M = (E, \mathcal{B})$ or in the form (M, r) where $\mathcal{B}$ is the set of its bases, and r is its rank function.

Lemma 17 *Let $M = (E, \mathcal{B})$ be a matroid, and $\mathcal{C}$ a set of cuts of M such that*

$$\mathcal{B}_{\mathcal{C}} := \{B \in \mathcal{B} : |B \cap C| \leq 1 \text{ for all } C \in \mathcal{C}\} \neq \emptyset.$$

Then $\mathcal{B}_{\mathcal{C}}$ is the set of bases of a matroid.

It follows immediately that then in fact $|B \cap C| = 1$ (for all $B \in \mathcal{B}_{\mathcal{C}}, C \in \mathcal{C}$); it follows from $\mathcal{B}_{\mathcal{C}} \neq \emptyset$ that each $C \in \mathcal{C}$ has elements that are not contained in any other member of $\mathcal{C}$. (Hint: If $b \in B \in \mathcal{B}$, then there exists a unique cut in M disjoint of $B \setminus \{b\}$.)

Proof: Let $B_1, B_2 \in \mathcal{B}_{\mathcal{C}}$, $b_2 \in B_2 \setminus B_1$ arbitrary. We show there exists $b_1 \in B_1$ such that $(B_1 \setminus \{b_1\}) \cup \{b_2\} \in \mathcal{B}_{\mathcal{C}}$.

If $b_2 \in \cup_{C \in \mathcal{C}} C$, then let C be the (unique) cut disjoint of $B_2 \setminus b_2$; then $C \in \mathcal{C}$. But then the unique circuit of $B_1 \cup \{b_2\}$ must have an element in C besides b_2, denote it b_1. (Since $B_1 \setminus \{b_1\}$ has also a unique disjoint cut, C is the only cut of $\mathcal{C}$ containing b_1.) So $(B_1 \setminus \{b_1\}) \cup \{b_2\} \in \mathcal{B}_{\mathcal{C}}$.

If $b_2 \notin \cup_{C \in \mathcal{C}} C$, then the unique circuit of $B_1 \cup \{b_2\}$ is disjoint of $\cup \mathcal{C}$, since a circuit and a cut cannot meet in exactly 1 element. Let $b_1 \neq b_2$ be an element of this circuit. □